

Chapter 1

Introduction

1.1 WHAT IS GAME THEORY?

We, humans, cannot survive without interacting with other humans, and ironically, it sometimes seems that we have survived despite those interactions. Production and exchange require cooperation between individuals at some level but the same interactions may also lead to disastrous confrontations. Human history is as much a history of fights and wars as it is a history of successful cooperation. Many human interactions carry the potentials of cooperation and harmony as well as conflict and disaster. Examples are abound: relationships among couples, siblings, countries, management and labor unions, neighbors, students and professors, and so on.

One can argue that the increasingly complex technologies, institutions, and cultural norms that have existed in human societies have been there in order to facilitate and regulate these interactions. For example, internet technology greatly facilitates buyer-seller transactions, but also complicates them further by increasing opportunities for cheating and fraud. Workers and managers have usually opposing interests when it comes to wages and working conditions, and labor unions as well as labor law provide channels and rules through which any potential conflict between them can be addressed. Similarly, several cultural and religious norms, such as altruism or reciprocity, bring some order to potentially dangerous interactions between individuals. All these norms and institutions constantly evolve as the nature of the underlying interactions keep changing. In this sense, understanding human behavior in its social and institutional context requires a proper understanding of human interaction.

Economics, sociology, psychology, and political science are all devoted to studying human behavior in different realms of social life. However, in many instances they treat individuals in isolation, for convenience if not for anything else. In other words, they assume that to understand

VISIT...

LANZAROTE
Caliente.COM

one individual's behavior it is safe to assume that her behavior does not have a significant effect on other individuals. In some cases, and depending upon the question one is asking, this assumption may be warranted. For example, what a small farmer in a local market, say in Montana, charges for wheat is not likely to have an effect on the world wheat prices. Similarly, the probability that my vote will change the outcome of the U.S. presidential elections is negligibly small. So, if we are interested in the world wheat price or the result of the presidential elections, we may safely assume that one individual acts as if her behavior will not affect the outcome.

In many cases, however, this assumption may lead to wrong conclusions. For example, how much our farmer in Montana charges, compared to the other farmers in Montana, certainly affects how much she and other farmers make. If our farmer sets a price that is lower than the prices set by the other farmers in the local market, she would sell more than the others, and vice versa. Therefore, if we assume that they determine their prices without taking this effect into account, we are not likely to get anywhere near understanding their behavior. Similarly, the vote of one individual may radically change the outcome of voting in small committees and assuming that they vote in ignorance of that fact is likely to be misleading.

The subject matter of game theory is exactly those interactions within a group of individuals (or governments, firms, etc.) where the actions of each individual have an effect on the outcome that is of interest to all. Yet, this is not enough for a situation to be a proper subject of game theory: the way that individuals act has to be strategic, i.e., they should be aware of the fact that their actions affect others. The fact that my actions have an effect on the outcome does not necessitate strategic behavior, if I am not aware of that fact. Therefore, we say that game theory studies *strategic interaction* within a group of individuals. By strategic interaction we mean that individuals know that their actions will have an effect on the outcome and act accordingly.

Having determined the types of situations that game theory deals with, we have to now discuss how it analyzes these situations. Like any other theory, the objective of game theory is to organize our knowledge and increase our understanding of the outside world. A scientific theory tries to abstract the most essential aspects of a given situation, analyze them using certain assumptions and procedures, and at the end derive some general principles and predictions that can be applied to individual instances.

For it to have any predictive power, game theory has to postulate some rules according to which individuals act. If we do not describe how individuals behave, what their objectives are and how they try to achieve those objectives we cannot derive any predictions at all in a given situation. For example, one would get completely different predictions regarding the price of wheat in a local market if one assumes that farmers simply flip a coin and choose between \$1 and \$2 a pound compared to if one assumes they try to make as much money as possible. Therefore, to bring some

discipline to the analysis one has to introduce some structure in terms of the rules of the game.

The most important, and maybe one of the most controversial, assumption of game theory which brings about this discipline is that individuals are *rational*.

We assume that individuals are rational.

Definition. An individual is *rational* if she has well-defined objectives (or preferences) over the set of possible outcomes and she implements the best available strategy to pursue them.

Rationality implies that individuals know the strategies available to each individual, have complete and consistent preferences over possible outcomes, and they are aware of those preferences. Furthermore, they can determine the best strategy for themselves and flawlessly implement it.

 If taken literally, the assumption of rationality is certainly an unrealistic one, and if applied to particular cases it may produce results that are at odds with reality. We should first note that game theorists are aware of the limitations imposed by this assumption and there is an active research area studying the implications of less demanding forms of rationality, called *bounded rationality*. This course, however, is not the appropriate place to study this area of research. Furthermore, to really appreciate the problems with rationality assumption one has to first see its results. Therefore, without delving into too much discussion, we will argue that one should treat rationality as a limiting case. You will have enough opportunity in this book to decide for yourself whether it produces useful and interesting results. As the saying goes: “the proof of the pudding is in the eating.”

The term strategic interaction is actually more loaded than it is alluded to above. It is not enough that I know that my actions, as well as yours, affect the outcome, but I must also know that you know this fact. Take the example of two wheat farmers. Suppose both farmer A and B know that their respective choices of prices will affect their profits for the day. But suppose, A does not know that B knows this. Now, from the perspective of farmer A, farmer B is completely ignorant of what is going on in the market and hence farmer B might set any price. This makes farmer A's decision quite uninteresting itself. To model the situation more realistically, we then have to assume that they both know that they know that their prices will affect their profits. One actually has to continue in this fashion and assume that the rules of the game, including how actions affect the participants and individuals' rationality, are common knowledge.

A fact X is *common knowledge* if everybody knows it, if everybody knows that everybody knows it, if everybody knows that everybody knows that everybody knows it, and so on. This has

We assume that the game and
rationality are common
knowledge

some philosophical implications and is subject to a lot of controversy, but for the most part we will avoid those discussions and take it as given.

In sum, we may define game theory as follows:

Definition. *Game theory* is a systematic study of strategic interaction among rational individuals.

Its limitations aside, game theory has been fruitfully applied to many situations in the realm of economics, political science, biology, law, etc. In the rest of this chapter we will illustrate the main ideas and concepts of game theory and some of its applications using simple examples. In later chapters we will analyze more realistic and complicated scenarios and discuss how game theory is applied in the real world. Among those applications are firm competition in oligopolistic markets, competition between political parties, auctions, bargaining, and repeated interaction between firms.

1.2 EXAMPLES

For the sake of comparison, we first start with an example in which there is no strategic interaction, and hence one does not need game theory to analyze.

Example 1.1 (A Single Person Decision Problem). Suppose Ali is an investor who can invest his \$100 either in a safe asset, say government bonds, which brings 10% return in one year, or he can invest it in a risky asset, say a stock issued by a corporation, which either brings 20% return (if the company performance is good) or zero return (if the company performance is bad).

	State	
	Good	Bad
Bonds	10%	10%
Stocks	20%	0%

Clearly, which investment is best for Ali depends on his preferences and the relative likelihoods of the two states of the world. Let's denote the probability of the good state occurring p and that of the bad state $1 - p$, and assume that Ali wants to maximize the amount of money he has at the end of the year. If he invests his \$100 on bonds, he will have \$110 at the end of the year irrespective of the state of the world (i.e., with certainty). If he invests on stocks, however, with probability

p he will have \$120 and with probability $1 - p$ he will have \$100. We can therefore calculate his average (or expected) money holdings at the end of the year as

$$p \times 120 + (1 - p) \times 100 = 100 + 20 \times p$$

If, for example, $p = 1/2$, then he expects to have \$110 at the end of the year. In general, if $p > 1/2$, then he would prefer to invest in stocks, and if $p < 1/2$ he would prefer bonds.

This is just one example of a *single person decision making problem*, in which the decision problem of an individual can be analyzed in isolation of the other individuals' behavior. Any uncertainty involved in such problems are exogenous in the sense that it is not determined or influenced in any way by the behavior of the individual in question. In the above example, the only uncertainty comes from the performance of the stock, which we may safely assume to be independent of Ali's choice of investment. Contrast this with the situation illustrated in the following example.

A single person decision problem has no strategic interaction

Example 1.2 (An Investment Game). Now, suppose Ali again has two options for investing his \$100. He may either invest it in bonds, which have a certain return of 10%, or he may invest it in a risky venture. This venture requires \$200 to be a success, in which case the return is 20%, i.e., \$100 investment yields \$120 at the end of the year. If total investment is less than \$200, then the venture is a failure and yields zero return, i.e., \$100 investment yields \$100. Ali knows that there is another person, let's call her Beril, who is exactly in the same situation, and there is no other potential investor in the venture. Unfortunately, Ali and Beril don't know each other and cannot communicate. Therefore, they both have to make the investment decision without knowing the decisions of each other.

We can summarize the returns on the investments of Ali and Beril as a function of their decisions in the table given in Figure 1.1. The first number in each cell represents the return on Ali's investment, whereas the second number represents Beril's return. We assume that both Ali and Beril know the situation represented in this table, i.e., they know the rules of the game.

Figure 1.1: Investment Game.

		Beril	
		Bonds	Venture
Ali	Bonds	110, 110	110, 100
	Venture	100, 110	120, 120

The existence of strategic interaction is apparent in this situation, which should be contrasted with the one in Example 1.1. The crucial element is that the outcome of Ali's decision (i.e., the return on the investment chosen) depends on what Beril does. Investing in the risky option, i.e., the

venture, has an uncertain return, as it was the case in Example 1.1. However, now the source of the uncertainty is another individual, namely Beril. If Ali believes that Beril is going to invest in the venture, then his optimal choice is the venture as well, whereas, if he thinks Beril is going to invest in bonds, his optimal choice is to invest in bonds. Furthermore, Beril is in a similar situation, and this fact makes the problem significantly different from the one in Example 1.1.

So, what should Ali do? What do you expect would happen in this situation? At this point we do not have enough information in our model to provide an answer. First we have to describe Ali and Beril's objectives, i.e., their preferences over the set of possible outcomes. One possibility, economists' favorite, is to assume that they are both expected payoff, or utility, maximizers. If we further take utility to be the amount of money they have, then we may assume that they are expected money maximizers. This, however, is not enough for us to answer Ali's question, for we have to give Ali a way to form expectations regarding Beril's behavior.

One simple possibility is to assume that Ali thinks Beril is going to choose bonds with some given probability p between zero and one. Then, his decision problem becomes identical to the one in Example 1.1. Under this assumption, we do not need game theory to solve his problem. But, is it reasonable for him to assume that Beril is going to decide in such a mechanical way? After all, we have just assumed that Beril is an expected money maximizer as well. So, let's assume that they are both rational, i.e., they choose whatever action that maximizes their expected returns, and they both know that the other is rational.

Is this enough? Well, Ali knows that Beril is rational, but this is still not enough for him to deduce what she will do. He knows that she will do what maximizes her expected return, which, in turn, depends on what she thinks Ali is going to do. Therefore, what Ali should do depends on what she thinks Beril thinks that he is going to do. So, we have to go one more step and assume that not only each knows that the other is rational but also each knows that the other knows that the other is rational. We can continue in this manner to argue that an intelligent solution to Ali's conundrum is to assume that both know that both are rational; both know that both know that both are rational; both know that both know that both know that both are rational; ad infinitum. This is a difficult problem indeed and game theory deals exactly with this kind of problems. The next example provides a problem that is relatively easier to solve.

Example 1.3 (Prisoners' Dilemma). Probably the best known example, which has also become a parable for many other situations, is called the Prisoners' Dilemma. The story goes as follows: two suspects are arrested and put into different cells before the trial. The district attorney, who is pretty sure that both of the suspects are guilty but lacks enough evidence, offers them the following deal: if both of them confess and implicate the other (labeled C), then each will be sentenced to, say, 5 years of prison time. If one confesses and the other does not (labeled N), then the "rat" goes

free for his cooperation with the authorities and the non-confessor is sentenced to 6 years of prison time. Finally, if neither of them confesses, then both suspects get to serve one year.

We can compactly represent this story as in Figure 1.2 where each number within each cell is the number of free years that will be enjoyed by each prisoner in the next six years.

Figure 1.2: Prisoners' Dilemma.

		Player 2	
		C	N
Player 1	C	-5, -5	0, -6
	N	-6, 0	-1, -1

For instance, the best outcome for the player 1 is the case in which he confesses and the player 2 does not. The next best outcome for player 1 is (N, N) , and then (C, C) and finally (N, C) . A similar interpretation applies to player 2.

How would you play this game in the place of player 1? One useful observation is the following: no matter what player 2 intends to do, playing C yields a better outcome for player 1. This is so because (C, C) is a better outcome for him than (N, C) , and (C, N) is a better outcome for him than (N, N) . So, it seems only “rational” for player 1 to play C by confessing. The same reasoning for player 2 entails that this player too is very likely to play C . A very reasonable prediction here is, therefore, that the game will end in the outcome (C, C) in which both players confess to their crimes.

And this is the dilemma: wouldn't each of the players be strictly better off by playing N instead? After all, (N, N) is preferred by both players to (C, C) . It is really a pity that the rational individualistic play leads to an inferior outcome from the perspective of both players.

You may at first think that this situation arises here only because the prisoners are put into separate cells and hence are not allowed to have pre-play communication. Surely, you may argue, if the players debate about how to play the game, they would realize that (N, N) is superior relative to (C, C) for both of them, and thus agree to play N instead of C . But even if such a verbal agreement is reached prior to the actual play of the game, what makes player 1 so sure that player 2 will not backstab him in the last instant by playing C ; after all, if player 2 is convinced that player 1 will keep his end of the bargain by playing N , it is better for her to play C . Thus, even if such an agreement is reached, both players may reasonably fear betrayal, and may thus choose to betray before being betrayed by playing C ; we are back to the dilemma.



What do you think would happen if players could sign binding contracts?

Even if you are convinced that there is a genuine dilemma here, you may be wondering why

we are making such a big deal out of a silly story. Well, first note that the “story” of the prisoners’ dilemma is really only a story. The dilemma presented above correspond to far more realistic scenarios. The upshot is that there are instances in which the interdependence between individuals who rationally follow their self-interest yields socially undesirable outcomes. Considering that one of the main claims of the neoclassical economics is that selfish pursuit of individual welfare yields efficient outcomes (the famous invisible hand), this observation is a very important one, and economists do take it very seriously. We find in prisoners’ dilemma a striking demonstration of the fact that the classical claim that “decentralized behavior implies efficiency” is not necessarily valid in environments with genuine room for strategic interaction.

 Prisoners’ dilemma type situations actually arise in many interesting scenarios, such as arms-races, price competition, dispute settlements with or without lawyers, etc. The common element in all these scenarios is that if everybody is cooperative a good outcome results, but nobody finds it in her self-interest to act cooperatively, and this leads to a less desirable outcome. As an example consider the pricing game in a local wheat market (depicted in Figure 1.3) where there are only two farmers and they can either set a low price (L) or a high price (H). The farmer who sets the lowest price captures the entire market, whereas if they set the same price they share the market equally.

Figure 1.3: Pricing Game.

		Farmer B	
		L	H
Farmer A	L	1, 1	4, 0
	H	0, 4	2, 2

This example paints a very grim picture of human interactions. Indeed, many times we observe cooperation rather than its complete failure. One important area of research in game theory is the analysis of environments, institutions, and norms, which actually sustain cooperation in the face of such seemingly hopeless situations as the prisoners’ dilemma.

Just to illustrate one such scenario, consider a repetition of the Prisoners’ Dilemma game. In a repeated interaction, each player has to take into account not only what is their payoff in each interaction but also how the outcome of each of these interactions influences the future ones. For example, each player may induce cooperation by the other player by adopting a strategy that punishes bad behavior and rewards good behavior. We will analyze such repeated interactions in Chapter 9.

Example 1.4 (Rebel Without a Cause). In the classic 1955 movie *Rebel Without a Cause*, Jim, played by James Dean, and Buzz compete for Judy, played by Natalie Wood. Buzz's gang members gather by a cliff that drops down to the Pacific Ocean. Jim and Buzz are to drive toward the cliff; the first person to jump from his car is declared the chicken whereas the last person to jump is a hero and captures Judy's heart. Each player has two strategies: jump before the other player (B) and after the other player (A). If they jump at the same time (B, B), they survive but lose Judy. If one jumps before and the other after, the latter survive and gets Judy, whereas the former gets to live, but without Judy. Finally, if both choose to jump after the other (A, A), they die an honorable death.

The situation can be represented as in Figure 1.4.

Figure 1.4: Game of Chicken.

		Buzz	
		B	A
Jim	B	2, 2	1, 3
	A	3, 1	0, 0


The likely outcome is not clear. If Jim thinks Buzz is going to jump before him, then he is better off waiting and jumping after. On the other hand, if he thinks Buzz is going to wait him out, he better jumps before: he is young and there will be other Judys. In the movie Buzz's leather jacket's sleeve is caught on the door handle of his car. He cannot jump, even though Jim jumps. Both cars and Buzz plunge over the cliff.¹

Game of chicken is also used as a parable of situations which are more interesting than the above story. There are dynamic versions of the game of chicken called the *war of attrition*. In a war of attrition game, two individuals are supposed to take an action and the choice is the timing of that action. Both players desire to be the last to take that action. For example, in the game of chicken, the action is to jump. Therefore, both players try to wait each other out, and the one who concedes first loses.

Example 1.5 (Entry Game). In all the examples up to here we assumed that the players either choose their strategies simultaneously or without knowing the choice of the other player. We model such situations by using what is known as *Strategic (or Normal) Form Games*.

In some situations, however, players observe at least some of the moves made by other players and therefore this is not an appropriate modeling choice. Take for example the *Entry Game* depicted in Figure 1.5. In this game Pepsi (P) first decides whether to enter a market currently monopolized

¹In real life, James Dean killed himself and injured two passengers while driving on a public highway at an estimated speed of 100 mph.

Figure 1.5: Entry Game**Table 1.1: Voters' Preferences**

voter 1	voter 2	voter 3
A	B	S
S	A	A
B	S	B

by Coke (C). After observing Pepsi's choice Coke decides whether to fight the entry (F) by, for example, price cuts and/or advertisement campaigns, or acquiesce (A).


Such games of sequential moves are modeled using what is known as *Extensive Form Games*, and can be represented by a game tree as we have done in Figure 1.5.

In this example, we assumed that Pepsi prefers entering only if Coke is going to acquiesce, and Coke prefers to stay as a monopoly, but if entry occurs it prefers to acquiesce; hence the payoff numbers appended to the end nodes of the game.

- ☞ What do you think Pepsi should do?
- ☞ Is there a way for Coke to avoid entry?

Example 1.6 (Voting). Another interesting application of game theory, to political science this time, is *voting*. As a simple example, suppose that there are two competing bills, A and B, and three legislators, voters 1, 2 and 3, who are to vote on these bills. The voting takes place in two stages. They first vote between A and B, and then between the winner of the first stage and the status-quo, denoted S. The voters' rankings of the alternatives are given in Table 1.1.

First note that if each voter votes truthfully, A will be the winner in the first round, and it will also win against the status-quo in the second round. Do you think this will be the outcome? Well, voter 3 is not very happy about the outcome and has another way to vote which would make him

Figure 1.6: Voting Game

happier. Assuming that the other voters keep voting truthfully, she can vote for B, rather than A, in the first round, which would make B the winner in the first round. B will lose to S in the second round and voter 3 is better off. Could this be the outcome? Well, now voter 2 can switch her vote to A to get A elected in the first round which then wins against S. Since she likes A better than S she would like to do that.

We can analyze the situation more systematically starting from the second round. In the second round, each voter should vote truthfully, they have nothing to gain and possibly something to lose by voting for a less preferred option. Therefore, if A is the winner of the first round, it will also win in the second round. If B wins in the first round, however, the outcome will be S. This means that, by voting between A and B in the first round they are actually voting between A and S. Therefore, voter 1 and 2 will vote for A and eventual outcome will be A. (see Figure 1.6.)

Example 1.7 (Investment Game with Incomplete Information). So far, in all the examples, we have assumed that every player knows everything about the game, including the preferences of the other players. Reality, however, is not that simple. In many situations we lack relevant information regarding many components of a strategic situation, such as the identity and preferences of other players, strategies available to us and to other players, etc. Such games are known as *Games with Incomplete (or Private) Information*.

As an illustration, let us go back to Example 1.2, which we modify by assuming that Ali is not certain about Beril's preferences. In particular, assume that he believes (with some probability p) that Beril has the preferences represented in Figure 1.1, and with probability $1 - p$ he believes Beril is a little crazy and has some inherent tendency to take risks, even if they are unreasonable from the perspective of a rational investor. We represent the new situation in Figure 1.7.

Figure 1.7: Investment Game with Incomplete Information

		Beril	
		Bonds	Venture
Ali	Bonds	110, 110	110, 100
	Venture	100, 110	120, 120

Normal (p)

		Beril	
		Bonds	Venture
Ali	Bonds	110, 110	110, 120
	Venture	100, 110	120, 120

Crazy ($1 - p$)

If Ali was sure that Beril was crazy, then his choice would be clear: he should choose to invest in the venture. How small should p be for the solution of this game to be both Ali and Beril, irrespective of her preferences, investing in the venture? Suppose that “normal” Beril chooses bonds and Ali believes this to be the case. Investing in bonds yields \$110 for Ali irrespective of what Beril does. Investing in the venture, however, has the following expected return for Ali

$$p \times 100 + (1 - p) \times 120 = 120 - 20p$$


which is bigger than \$110 if $p < 1/2$. In other words, we would expect the solution to be investment in the venture for both players if Ali’s belief that Beril is crazy is strong enough.

Example 1.8 (Signalling). In Example 1.7 one of the players had incomplete information but they chose their strategies without observing the choices of the other player. In other words, players did not have a chance to observe others’ behavior and possibly learn from them. In certain strategic interactions this is not the case. When you apply for a job, for example, the employer is not exactly sure of your qualities. So, you try to impress your prospective boss with your resume, education, dress, manners etc. In essence, you try to *signal* your good qualities, and hide the bad ones, with your behavior. The employer, on the other hand, has to figure out which signals she should take seriously and which ones to discount (i.e. she tries to *screen* good candidates).

This is also the case when you go out on a date with someone for the first time. Each person tries to convey their good sides while trying to hide the bad ones, unless of course, it was a failure from the very beginning. So, there is a complex interaction of signalling and screening going on. Suppose, for example, that Ali takes Beril out on a date. Beril is going to decide whether she is going to have a long term relationship with him (call that marrying) or dump him. However, she wants to marry a smart guy and does not know whether Ali is smart or not. However, she thinks he is smart or dumb with equal probabilities. Ali really wants to marry her and tries to show that he is smart by cracking jokes and being funny in general during the date. However, being funny is not very easy. It is just stressful, and particularly so if one is dumb, to constantly try to come up with jokes that will impress her. Figure 1.8 illustrates the situation.

What do you think will happen at the end? Is it possible for a dumb version of Ali to be funny and marry Beril? Or, do you think it is more likely for a smart Ali to marry Beril by being funny, while a dumb Ali prefers to be quite and just enjoys the food, even if the date is not going further than the dinner?

Example 1.9 (Hostile Takeovers). During the 1980s there was a huge wave of mergers and acquisitions in the United States. Many of the acquisitions took the form of “hostile takeovers,” a term used to describe takeovers that are implemented against the will of the target company’s manage-

Figure 1.8: Dating Game

ment. They usually take the form of direct tender offers to shareholders, i.e., the acquirer publicly offers a price to all the shareholders. Some of these tender offers were in the form of what is known as “two-tiered tender offer.”

Such was the case in 1988 when Robert Campeau made a tender offer for Federated Department Stores. Let us consider a simplified version of the actual story. Suppose that the pre-takeover price of a Federated share is \$100. Campeau offers to pay \$105 per share for the first 50% of the shares, and \$90 for the remainder. All shares, however, are bought at the average price of the total shares tendered. If the takeover succeeds, the shares that were not rendered are worth \$90 each.

For example, if 75% of the shares are tendered, Campeau pays \$105 to the first 50% and pays \$90 to the remaining 25%. The average price that Campeau pays is then equal to

$$\begin{aligned}
 p &= 105 \times \frac{50}{75} + 90 \times \frac{25}{75} \\
 &= 100
 \end{aligned}$$

In general, if s percent of the shares are tendered the average price paid by Campeau, and thus

the price of a tendered share, is given by

$$p = \begin{cases} 105 & \text{if } s \leq 50 \\ 105 \times \frac{50}{s} + 90 \times \frac{s-50}{s} & \text{if } s > 50 \end{cases}$$

Notice that if everybody tenders, i.e., $s = 100$, then Campeau pays \$97.5 per share which is less than the current market price. So, this looks like a good deal for Campeau, but only if sufficiently high number of shareholders tender.

- 👉 If you were a Federated shareholder, would you tender your shares to Campeau?
- 👉 Does your answer depend on what you think other shareholders will do?
- 👉 Now suppose Macy's offers \$102 per share conditional upon obtaining the majority. What would you do?

The actual unfolding of events were quite unfortunate for Campeau. Macy's joined the bidding and this increased the premium quite significantly. Campeau finally won out (not by a two-tiered tender offer, however) but paid \$8.17 billion for the stock of a company with a pre-acquisition market value of \$2.93 billion. Campeau financed 97 percent of the purchase price with debt. Less than two years later, Federated filed for bankruptcy and Campeau lost his job.

1.3 OUR METHODOLOGY

So, we have seen that many interesting situations involve strategic interactions between individuals and therefore render themselves to a game theoretical study. At this point one has two options. We can either analyze each case separately or we may try to find general principals that apply to any game. As we have mentioned before, game theory provides tools to analyze strategic interactions, which may then be applied to any arbitrary game-like situation. In other words, throughout this course we will analyze abstract games, and suggest "reasonable" outcomes as solutions to those games. To fix ideas, however, we will discuss applications of these abstract concepts to particular cases which we hope you will find interesting.

We will analyze games along two different dimensions: (1) the order of moves; (2) information. This gives us four general forms of games, as we illustrate in Table 1.2.

Table 1.2: Game Forms

		Information	
		Complete	Incomplete
Moves	Simultaneous	Strategic Form Games with Complete Information <i>Example 1.2</i>	Bayesian Games <i>Example 1.7</i>
	Sequential	Extensive form Games with Complete Information <i>Example 1.5</i>	Extensive form Games with Incomplete Information <i>Example 1.8</i>

Strategic Form Games with Complete Information

Levent Koçkesen, Columbia University

Efe A. Ok, New York University

The simplest form of strategic interdependence prevails in contexts in which the actions taken simultaneously by a number of individuals determine the outcome of the interaction at once. To model such a setting, all we need to do is to specify the set of interacting individuals (commonly called the players), the set of actions available to these individuals, and a description of the incentives regarding the modeled interaction. That is, we need to write down the who, what and why of the setting we are trying to model.

Formally speaking, we need exactly three objects to define a **game in strategic form**. First we need to specify

- *the set of players*: $N = \{1, \dots, n\}$.

Of course, one may choose to name the players by other means than using integers. Any choice is as arbitrary as the other in this regard.

Next we need to specify

- *the action (strategy) spaces of players*: $A_i, i = 1, \dots, n$.

We interpret A_i as the set of all available actions to player i . That is, for player i , “playing the game” means choosing an action from the set A_i . For instance, in the children game “rock, scissors and paper,” the action space for each player is {rock, scissors, paper}, and in the prisoners’ dilemma the action space for each player i is {confess, not confess} (i.e. {C,N}).

Given the action spaces of the players, we define the *outcome space* of the game as

$$A = \times_{i \in N} A_i = \{(a_1, \dots, a_n) : a_i \in A_i, i = 1, \dots, n\}.$$

An *outcome* $a = (a_1, \dots, a_n)$ is thus nothing but an action profile.

To be able to formulate the decision problem of a player in a given strategic environment, we need to know about the preferences of this individual. To this end, we introduce

- *the payoff functions of players*: $u_i : A \rightarrow \mathbf{R}, i = 1, \dots, n$.

The interpretation of a payoff function is identical to that of a utility function that you have surely encountered in a microeconomics course. If $u_i(a_1, \dots, a_n) > u_i(b_1, \dots, b_n)$, then we understand that player i likes outcome $a = (a_1, \dots, a_n)$ strictly better than the outcome $b = (b_1, \dots, b_n)$. The crucial observation is that the payoff of the player i depends on not only the action chosen by player i but also on the action choices of the rest of the participating

players. As we have discussed before, this is a crucial element distinguishing a game theoretic decision problem from a single agent decision problem.

We should note that, at this level of generality, we treat a statement like $u_i(a) > u_i(b)$ as purely *ordinal*, that is, without attaching any meaning to the difference $u_i(a) - u_i(b)$. All we know in the formulation so far is how the individuals *rank* the outcomes, not how much “utils” they derive from them. For instance, in the prisoners’ dilemma game considered above, assigning payoff 0 for player 1 to the outcome (N,C) was arbitrary; any number would do so long as it is strictly smaller than that assigned to (C,C) (which must be strictly smaller than that assigned to (N,N) which, in turn, must be strictly smaller than that assigned to (C,N)). In later sections, when we start analyzing cases in which individuals face uncertainty, we will have to modify this assumption.

Summing up, we define formally a game in strategic form as the tuple

$$(N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N}).$$

(Note that the term “normal-form game” is also used in the literature.) Thus, when we talk about a “game in strategic form” we have in mind a setup in which all this information is provided. In particular, if the game is played by only two players (so that $N = \{1, 2\}$), we need exactly four pieces of information:

$$(A_1, A_2, u_1, u_2).$$

Therefore, if each player has finitely many actions available to him/her, then we can represent a 2-person game in strategic form by means of a bimatrix, as we have done in the previous section for the prisoners’ dilemma. In such a representation our convention is always that player 1 (who is a male) chooses the rows and player 2 (who is a female) chooses the columns.

The conceptual point to make is that a game in strategic form attempts to capture a scenario in which there is strategic interdependence among a set of players who either take their actions simultaneously or *without* observing the actions chosen by others. Therefore, strategic form games are not suitable to model situations where individuals take actions in a sequential manner after observing the actions taken before.

Remark. A crucial assumption behind a game in strategic form is that everything about the formulation of the game (that is the set of players, the set of actions and the utility functions) are all known by each player in the game. What is more, each player knows that all players know everything about the game, and all players know that each player knows that all players know everything about the game, and so on. Believe it or not, at a philosophical level, all this matters. But we shall not concern ourselves much with this issue; we shall simply postulate that the primitives of a game is *common knowledge* without worrying too

much about what this really means.¹ What is more, there is no uncertainty pertaining to the actions available to the players and to their payoff functions. This makes the game form defined in this section a strategic form game with complete information. In later sections we will have a chance to see how to model situations involving dynamic interaction as well as incomplete information on the part of some players. $\square$

As with any other new concept, the best way to come into grips with the games in strategic form is to study several specific examples, hence the next section.

¹Trying to model the idea of “common knowledge” requires some mathematical sophistication which is best avoided at this stage. We shall thus do no more on this topic than recommending to the interested reader the excellent survey of J. Geanakoplos (1992), “Common knowledge,” *Journal of Economic Perspectives* 6, pp. 53-82.

Examples

(1) *Prisoners' Dilemma. (PD)* Recall that the prisoners' dilemma scenario we have discussed in the introduction was represented by the bimatrix

	C	N
C	1,1	6,0
N	0,6	5,5

Here we have $A_1 = \{C,N\} = A_2$ so that $A = \{(C,C), (C,N), (N,C), (N,N)\}$, and $u_1(C,C) = 1$, $u_2(N,C) = 6$, ... You should make sure you understand that the bimatrix

	C	N
C	β, β	θ, α
N	α, θ	γ, γ

also represents the same game whenever $\alpha < \beta < \gamma < \theta$.

(2) *Battle of the Sexes. (BoS)* Consider the following scenario: Two individuals are arguing about what to do for entertainment in the evening. Player 1 wishes to go to a movie while player 2 wishes to go to an opera. (According to the somewhat sexist version of the story, player 1 is a man and player 2 is a woman; hence the name of the game.) However, the couple is in love, so, the most important thing for both players is to do something together; both view the night “wasted” unless they spend it together.

We may represent this story as a 2-person game in strategic form by means of the following *bimatrix*

	m	o
m	2,1	0,0
o	0,0	1,2

Here we have $A_1 = \{m,o\} = A_2$ so that $A = \{(m,m), (m,o), (o,m), (o,o)\}$, and $u_1(o,m) = 0$, $u_2(o,o) = 2$, ... (Once again the choice of utility values is arbitrary other than the ranking of the outcomes it entails.) Like the prisoners' dilemma, the BoS is also a famous example in game theory that will help us illustrate many interesting concepts later on. So perhaps now is a good time for you think about how you would play this game in actuality. What is your prediction about the outcome of this game without preplay communication? With preplay communication?

(3) *A Pure Coordination Game. (CG)* Consider two individuals each of whom chooses either l (for left) or r (for right). If their choices do not match, they receive a payoff of zero. On the other hand, if they coordinate on l, each receives a payoff of 1, and if they coordinate

on r, then each receives a payoff of 2. So, we have $A_1 = A_2 = \{l,r\}$ and the game being played is represented by the bimatrix

	l	r
l	1,1	0,0
r	0,0	2,2

This game too is an interesting one, and we shall come back to it later when we discuss the effects of preplay communication among the players. For now, ask yourself if your prediction about how this game would actually be played depends on whether preplay communication is allowed or not.

(4) *Merchantship-Warship Game (MW)* Consider the following scenario: two ships are approaching to an island, one from the east, one from the west. Both ships need to pass by the island so they are contemplating about whether they should do so by moving north (N) or south (S). The trick is that both ships know that the other ship is also approaching to the island, and while one of the them (player 1, a merchantship) wishes to avoid the other ship (player 2, an enemy warship), the warship wishes to come across the merchantship within the vicinity of the island.

This story is compactly represented as a 2-person game in strategic form by means of the following bimatrix:

	N	S
N	0,1	1,0
S	1,0	0,1

As an easy exercise, write down the outcome space of this game algebraically, and provide another bimatrix that corresponds to the same scenario described above. Do you see an “obvious” way of playing this game.

Before considering an economically motivated example, let us note that the PD and CG are (two-person) *symmetric games* in the sense that they satisfy the following two conditions:

- (i) $A_1 = A_2$, that is, each player has the same action set
- (ii) $u_1(a_1, a_2) = u_2(a_2, a_1)$ for all $a_i \in A_i$, $i = 1, 2$, that is, if we exchange the actions of the players (or in the n -person case, any two players holding everyone else’s action fixed), then the payoffs of these players are also exchanged. (Notice that this implies that the two players must receive the same payoff when they both choose the same action.)

For instance, in PD, $u_1(C,C) = 1 = u_2(C,C)$, $u_1(C,N) = 6 = u_2(N,C)$, etc. On the other hand, the BoS and MW are not symmetric. Here are two other asymmetric games.

	N	S			L	M	R
N	0,1	2,0	and	U	1,0	1,2	0,1
S	1,0	0,1			D	0,3	0,1

The symmetric games are in general simpler than asymmetric games because reasoning from the point of view of one player is sufficient in such games to understand how the other players reason as well. We shall utilize this fact in many examples that we shall consider in this book.

Let us now examine a slightly more sophisticated example of a game in strategic form. This example plays a fundamental role in the theory of industrial organization, and we shall work out several variations of it in the sequel.

(5) *The Cournot Duopoly Model*

Consider a market for a single (homogeneous) good whose market inverse demand function is

$$P = D(Q), \quad Q \geq 0$$

where P is the price of the good and Q is the quantity demanded. We assume that the function D is monotonically decreasing. Suppose that there are exactly two firms producing this good. The cost functions of these firms are

$$C_i = C_i(Q_i), \quad Q_i \geq 0, i = 1, 2,$$

where C_i is a twice differentiable function defined on $\mathbf{R}_+$ with $C'_i > 0$ and $C''_i \leq 0$.

We may model the market interaction of these firms as a 2-person game in strategic form as follows:

(i) $N = \{1, 2\}$.

(ii) $A_i = [0, \bar{Q}]$; thus $(Q_1, Q_2) \in A = [0, \bar{Q}]^2$ means that firm i is producing Q_i units at the outcome (Q_1, Q_2) . The value $\bar{Q} > 0$ is an upper bound on the level of production of firms acting as a *capacity constraint*.

(iii) $u_i(Q_1, Q_2) = D(Q_1 + Q_2)Q_i - C_i(Q_i)$ for each $Q_i \in A_i$, $i = 1, 2$.

This model is mathematically too general to allow for a relatively easy analysis. For this reason, a common specification adopted in the literature posits in addition that the firms operate under identical *constant unit costs* so that $C_i(Q_i) = cQ_i$ and that the market demand is given as

$$P = \begin{cases} a - bQ, & 0 \leq Q \leq a/b \\ 0, & a/b < Q \end{cases}$$

with $a > c > 0$ and $b > 0$ being given parameters. To simplify the analysis further, we set $\bar{Q}$ in this specification equal to a/b ; this is meaningful since no firm would realistically produce an output level that exceeds a/b in this setting as this would entail making negative profits. We refer to this model as the *linear Cournot model*, and observe that the payoff function of firm i in this model is:

$$u_i(Q_1, Q_2) = \begin{cases} (a - b(Q_1 + Q_2))Q_i - cQ_i, & 0 \leq Q_1 + Q_2 \leq a/b \\ -cQ_i, & a/b < Q_1 + Q_2 \end{cases}$$

for each $Q_i \in A_i$, $i = 1, 2$. Therefore, the associated 2-person game in strategic form is symmetric (while this is not necessarily the case in the general model).

The important thing to note in the Cournot model is that, unlike the market structures of perfect competition (where all firms disregard the actions of other firms since each firm is assumed to be negligible in the market) and of monopoly (where there is no other firm around to matter), one firm's action does not alone determine the outcome. Thus, we need the apparatus of game theory to provide a prediction with respect to the market outcome.

We shall later on encounter many more examples of games in strategic form. But now it is time that we turn to the question of how to play a strategic game.

Dominant Strategy Equilibrium

The problem of a player in a strategic game is to decide upon an action to take without knowing which actions will be taken by her opponents. Therefore, each individual has to form a conjecture regarding the action choices of the other players, and this is not always an easy task. But, in some cases, this difficulty does not really arise, because there is an optimal way of taking an action *independently* of the intended play of the others. We have in fact already encountered such a situation in the prisoners' dilemma. Indeed, taking the noncooperative action of confessing C is optimal for, say player 1, in the prisoners' dilemma no matter what player 2 is planning to do. In this sense, we say that there is an “obvious” way of playing the PD for player 1 (and similarly for player 2): choosing C. We formalize such sure-fire actions in general as follows.

Notation. Let $A = \times_{j \in N} A_j$ be the outcome space of an n -person game in strategic form, and let $a = (a_1, \dots, a_n) \in A$. For each i , we let a_{-i} stand for the vector

$$(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in \times_{j \in N \setminus \{i\}} A_j$$

and write $a = (a_i, a_{-i})$. Clearly, a_{-i} is nothing but a profile of actions taken by all players in the game other than i . We denote the set of all such profiles conveniently as A_{-i} . Formally speaking, we have $A_{-i} = \times_{j \in N \setminus \{i\}} A_j$.

Definition. Given a game in strategic form, an action $a_i \in A_i$ **weakly dominates** action $b_i \in A_i$ for player i if

$$u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}$$

and

$$u_i(a_i, a_{-i}) > u_i(b_i, a_{-i}) \quad \text{for some } a_{-i} \in A_{-i}.$$

It **strictly dominates** b_i if

$$u_i(a_i, a_{-i}) > u_i(b_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}.$$

Definition. Given a game in strategic form, an action $a_i \in A_i$ is **weakly dominant** if it weakly dominates every action in A_i . It is called **strictly dominant** if it strictly dominates every action in A_i .

Remark. A dominant action must be unique. Why? $\square$

To reiterate, a dominant strategy for a player is an action that is optimal for this player no matter what his opponents do. Put differently, a player with a dominant action does not have to worry about how his opponents will play the game; for any belief that he might have about

the plans of actions by others, playing a dominant action is optimal. Consequently, there is good reason to believe that rational players would play their dominant actions in a given game (of course, provided that such actions are present). This idea leads us to the following equilibrium concept.

Definition. The **weakly dominant strategy equilibrium** of a game G in strategic form is defined as the weakly dominant action profile, and is denoted by $\mathbf{D}^w(G)$. Replacing the word “weakly” with “strictly” yields the definition for the *strictly* dominant strategy equilibrium, which is denoted by $\mathbf{D}^s(G)$.

Example. As we noted earlier, the action N is strictly dominant for both players in PD. Thus $\mathbf{D}^s(\text{PD}) = \{(C,C)\}$, which is also the weakly dominant strategy equilibrium as a strictly dominant action is also a weakly dominant action. As an example of a weakly dominant strategy equilibrium which is not strict, consider

	L	R
U	2,1	0,2
D	2,3	4,3

In this game there is no strictly dominant strategy equilibrium. There is, however, a weakly dominant strategy equilibrium given by the action profile (D,R).□

Dominant strategy equilibrium is quite a reasonable equilibrium concept which does not demand an excessive amount of “rationality” from the players. It only demands the players to be (rational) optimizers, and does not require them to know that the others are rational too. Unfortunately, this concept is silent in many interesting games since the existence of a dominant action for all players in a given game is a relatively rare phenomenon. (Check to see that there is no dominant strategy equilibrium in BoS, CG, MW, and Cournot oligopoly games.) It seems that we need to demand more rationality from the players to obtain more powerful predictions. We now turn to a systematic way of doing this.

Dominance Solvability

We have argued above that a “rational” player would play a dominant action (when such an action exists). Turning this argument on its head, we may then say that a “rational” player would never play an action when there is another action available to her that guarantees strictly more payoffs for this player irrespective of the intended play of others. We refer to such an action as a strictly dominated action. Formally,

Definition. Take a game in strategic form and consider any two actions $a_i, b_i \in A_i$ for any player $i \in N$. We say that a_i is **strictly dominated** by b_i if

$$u_i(a_i, a_{-i}) < u_i(b_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}.$$

We say that a_i is **weakly dominated** by b_i if

$$u_i(a_i, a_{-i}) \leq u_i(b_i, a_{-i}) \quad \text{for all } a_{-i} \in A_{-i}$$

while

$$u_i(a_i, a_{-i}) < u_i(b_i, a_{-i}) \quad \text{for some } a_{-i} \in A_{-i}.$$

A fundamental premise in game theory is that “rational” players do not play strictly dominated actions. For, as the argument goes, there is no belief that a player may hold about the intended play of others such that a strictly dominated action is optimal. Therefore, given a game G in strategic form, it makes sense to eliminate all the strictly dominated actions for any one of the players; after all “rational” players know that this player will not take any such action. But if all players ponder about how to play the game after eliminating (in their heads) strictly dominated actions of a given player, then the actual game being played is *effectively* a smaller game than the original one. But then why don’t we search for strictly dominated actions in this smaller game, that is, eliminate next the strictly dominated actions of another player *requiring “dominance” only against actions not yet eliminated*. And why not continue this way as far as we can?

Well, doing this may or may not be a reasonable thing to do depending on the context. Nevertheless, this elimination process, which is called the **iterated elimination of strictly dominated (IESD)** actions, certainly leads us to an interesting equilibrium concept. First of all, it yields an extension of the strictly dominant strategy equilibrium. While this is formally obvious, it is an important observation and we state it as a proposition (Proposition A below).

Moreover, the IESD actions may apply in many games with no dominant strategy equilibrium, and may yield a prediction concerning the play of the game even if no player has a dominant action. This prediction may even be sharp enough to entail a unique outcome. In this case we say that G is **dominance solvable**.

Proposition A. *If both players have strictly dominant actions, then IESD actions leads to the unique dominant strategy equilibrium; so such a game is dominance solvable.*

Proof. Obvious. ■

Example. The PD is dominance solvable by Proposition A. On the other hand, the IESD actions does not at all refine the outcome space in the BoS since neither of the players has a strictly dominated action in this game: BoS is not dominance solvable. To give a less trivial example, consider the game

	L	M	R
U	1,0	1,2	0,1
D	0,3	0,1	2,0

which does not possess a dominant strategy equilibrium. Observe that R is strictly dominated for player 2 (by action M). Therefore, in the first stage of the IESD process, we eliminate R. The idea is that player 1, being “rational,” knows that player 2 will not play R, and views the game effectively as

	L	M
U	1,0	1,2
D	0,3	0,1

But player 2, being “rational,” knows that player 1 is really contemplating about how to play this smaller game, and notices that in this game D is strictly dominated for player 1. So player 2 eliminates (in his head) the action D for player 1. This is the second stage of the IESD process and leaves us with the game

	L	M
U	1,0	1,2

We now reach to the final stage of the IESD process where we eliminate L for player 2. Hence this game is dominance solvable, and IESD actions leads to the outcome (U,M). □

Observe that applying the process of IESD actions in the case of a finite game is technically easy. In the example above, for instance, the outcome is immediately obtained by eliminating first R, then D and then L. However, you should keep in mind that the longer this process takes, the more “he knows that she knows that he knows that ...” sort of reasonings are used, and hence the “more rational” we demand the players should be. Put differently, for the IESD actions to make conceptual sense, not only that each player must not take strictly dominated actions, but also that each player must know that her opponent won’t do so, that he knows that her opponent knows that he won’t do so, and so on. So, this concept is less plausible in complicated games. Here is an example of a dominance solvable game which requires the

players to be, in a certain sense, “infinitely rational.” You should decide for yourself how reasonable is the IESD actions in this example.

Example. We consider the linear Cournot model. Observe that

$$\frac{du_1(Q_1, Q_2)}{dQ_1} = \begin{cases} (a - c) - 2bQ_1 - bQ_2, & 0 \leq Q_1 + Q_2 \leq a/b \\ -c, & a/b < Q_1 + Q_2 \end{cases}$$

so no matter what Q_2 is, $\frac{du_1}{dQ_1} < 0$ when $Q_1 > \frac{a-c}{2b}$ (This is perhaps a bit too swift, make sure you understand this step.) Thus any production level $Q_1 > \frac{a-c}{2b}$ is strictly dominated (by $\frac{a-c}{2b}$). In the first stage of the IESD actions process, therefore, we eliminate all $Q_i > \frac{a-c}{2b}$, $i = 1, 2$. Consequently, we have $Q_1 + Q_2 < a/b$ after one iteration. (As we discuss at the end of this section, the order of elimination does not matter for the final outcome in the case of IESD actions, so we can eliminate the strictly dominated actions of the firms simultaneously.) Consequently, given that $Q_2 \leq \frac{a-c}{2b}$, we have

$$\begin{aligned} \frac{du_1}{dQ_1} &= (a - c) - 2bQ_1 - bQ_2 \\ &\geq (a - c) - 2bQ_1 - b\left(\frac{a - c}{2b}\right) \\ &= \frac{a - c}{2} - 2bQ_1 \end{aligned}$$

so that $\frac{du_1}{dQ_1} > 0$ when $Q_1 < \frac{a-c}{4b}$. Thus, we eliminate all $Q_i < \frac{a-c}{4b}$, $i = 1, 2$. But, given that $Q_2 \geq \frac{a-c}{4b}$, one can similarly show that $\frac{du_1}{dQ_1} < 0$ when $Q_1 > \frac{3(a-c)}{8b}$. Iterating infinitely many times, then, only the outcome $(\frac{a-c}{3b}, \frac{a-c}{3b})$ survives the IESD actions. (*Challenge:* prove this.) Hence the linear Cournot model is dominance solvable. $\square$

We define the **iterated elimination of weakly dominated (IEWD)** actions in a way analogous to the IESD actions. But, as we shall see, this is a somewhat more problematic notion than IESD actions. To begin with there is a possible contradiction in the procedure: the argument behind not using weakly dominated actions is that if there is uncertainty in the mind of players as to the action choice of the other players then a weakly dominated action should not be used. For example, in the following game

	L	R
U	2,1	0,2
D	2,3	4,3

player 1 should not use action U if there is a small probability in his mind that player 2 will play R. Yet, the procedure of eliminating weakly dominated actions may involve deleting actions

to which the player previously assigned a positive probability. Consider in the following game

	L	R
U	3,1	2,0
M	4,0	1,1
D	4,4	2,4

we may delete U assuming that player 1 assigns a positive (however small) probability to the event that player 2 will play action L. Given U is deleted, then player 2's action L becomes weakly dominated and hence it can be deleted, i.e., it is going to be played with zero probability, contradicting the reason why action U was deleted to begin with.

Nevertheless, IEWD actions is used widely in economic applications of game theory, and we too will utilize this concept on occasion. Let us illustrate by means of two examples the power (and perhaps also the potential counter intuitiveness, you decide for yourself) of the notion of IEWD actions.

Example. (*Guess-the-average game*) Consider an n -person strategic game in which each player picks an integer between 1 and 999. So $N = \{1, \dots, n\}$ and $A_i = \{1, \dots, 999\}$. Let us write $\bar{a}$ for the mean of the action profile $(a_1, \dots, a_n)$, that is, $\bar{a} = \sum_{i=1}^n a_i/n$. The winners in this game are those players whose choice of integer is closest to $\frac{2}{3}\bar{a}$.²

- First take about five minutes to decide how you would play this game.
- Observe next that IESD actions does not provide a sharp prediction here; this game is not dominance solvable.
- Let us now apply IEWD actions. Take any player. This player knows that no matter what the other players play, the two-thirds of the average ballot cannot exceed 666. But then any integer larger than 666 is weakly dominated by 666 for this individual (why weakly?). Since this is true for all players, IEWD actions demands that we eliminate all actions in $\{667, \dots, 999\}$. But the argument can be repeated, for every strategy in $\{445, \dots, 666\}$ is *now* weakly dominated by 444. Continuing this way (iterating finitely many times), we find that the only outcome that survives the IEWD actions is $(1, \dots, 1)$!

□

²Formally, we may write $u_i(a_1, \dots, a_n) = 1$ if

$$\left| a_i - \frac{2}{3}\bar{a} \right| \leq \left| a_j - \frac{2}{3}\bar{a} \right| \text{ for all } j = 1, \dots, n$$

and $u_i(a_1, \dots, a_n) = 0$, otherwise.

Example. (*Chairman's Paradox*) Consider a committee of three persons (named as usual 1, 2 and 3) whose task is to choose an alternative from the choice set $\{\alpha, \beta, \gamma\}$ by means of voting. The alternative will be chosen on the basis of majority so that an alternative which gets two votes wins the election. The rule is such that if there is a tie (that is, if each voter votes for a different alternative, then the chairman of the committee, who is, say, player 3, will unilaterally decide on the outcome of the election by declaring the alternative that he best likes as the winner of the election. So this is not a symmetric game, it appears that the position of player 3 is strategically superior to the rest of the players.

Now assume that the preferences of the players are given as in the following list:

Player 1	Player 2	Player 3
α	β	γ
β	γ	α
γ	α	β

Here the convention is that any alternative in each column is strictly preferred to the alternatives that are below it by the corresponding player. For instance, player 1 strictly prefers α to β while she likes β strictly better than γ . Therefore, given these preferences, if all voters voted *sincerely*, each would vote for a different alternative, and in this case, player 3 would exert his additional power to declare the alternative γ as the winner of the election.

However, there is no reason why all voters should vote truthfully, in principal they would do so only if this would benefit them. What if they wish to play this voting game *strategically*? To see what would happen in this case, let us model the scenario as a game in strategic form where $A_i = \{\alpha, \beta, \gamma\}$ (an action for each individual is the vote that she is going to cast), and consider the following trimatrix representation:

	α	β	γ	
α	2,0,1	0,1,2	0,1,2	if player 3 chooses γ
β	0,1,2	1,2,0	0,1,2	
γ	0,1,2	0,1,2	0,1,2	
	α	β	γ	
α	2,0,1	2,0,1	2,0,1	if player 3 chooses α
β	2,0,1	1,2,0	0,1,2	
γ	2,0,1	0,1,2	0,1,2	
	α	β	γ	
α	2,0,1	1,2,0	0,1,2	if player 3 chooses β
β	1,2,0	1,2,0	1,2,0	
γ	2,0,1	1,2,0	0,1,2	

Here, for instance, $u_3(\alpha, \beta, \gamma) = 2$ since in this case the outcome of the election is γ which is the most preferred outcome by player 3. (Check that this representation really corresponds to the scenario described above.)

Our task is now to apply the IEWD actions to this game. Here is one way of doing this: Eliminate (1) γ for player 1; (2) α and γ for player 2; (3) α and β for player 3; (4) α for player 1. Hence the IEWD actions leads to the outcome (β, β, γ) which means that the winner of the election is β . Observe that this outcome contrasts sharply with the outcome in the case of sincere voting. In fact, with strategic voting, we observe that the worst outcome is elected for player 3 (if you believe in IEWD actions) who supposedly is a more powerful player than the others; this is why the present game is sometimes called *the chairman's paradox*. (What do you think is the key to “explain” this paradoxical outcome? What if players did not know the preferences of the others? What if they didn't believe that the others were so terribly smart? Do you agree with the prediction reached through the IEWD actions?) $\square$

Remark. An important question that we have to deal with before we conclude this section is this: could eliminating IESD actions lead to different results if elimination takes place in different orders? Fortunately, the answer is no. (Can you prove this?) However, the answer would be yes if we rather used weakly dominated actions in the iterations. For instance, consider the 2-person game in strategic form given by the bimatrix

	L	R
U	3,1	2,0
M	4,0	1,1
D	4,4	2,4

Here if we first eliminate player 1's action U, then player 2's action L, and then player 1's action M we get the outcome (D,R), while if we first eliminate M (and then R and then U) we get the outcome (D,L). This shows that the order of elimination matters in the case of IEWD actions. $\square$

The Nash Equilibrium

Levent Koçkesen, Columbia University

Efe A. Ok, New York University

As we have mentioned in our first lecture, one of the assumptions that we will maintain throughout is that individuals are rational, i.e., they take the best available actions to pursue their objectives. This is not any different from the assumption of rationality, or optimizing behavior, that you must have come across in your microeconomics classes. In most of microeconomics, individual decision making boils down to solving the following problem:

$$\max_{x \in X} u(x, \theta)$$

where x is the vector of choice variables, or possible actions, (such as a consumption bundle) of the individual, X denotes the set of possible actions available (such as the budget set), θ denotes a vector of parameters that are outside the control of the individual (such as the price vector and income), and u is the utility (or payoff) function of the individual.

What makes a situation a strategic game, however, is the fact that what is best for one individual, in general, depends upon other individuals' actions. The decision problem of an individual in a game can still be phrased in above terms by treating θ as the choices of other individuals whose actions affect the subject individual's payoff. In other words, letting $x = a_i$, $X = A_i$, and $\theta = a_{-i}$, the decision making problem of player i in a game becomes

$$\max_{a_i \in A_i} u_i(a_i, a_{-i}).$$

The main difficulty with this problem is the fact that an individual does not, in general, know the action choices of other players, a_{-i} , whereas in single decision making problems the parameter vector, θ , is assumed to be known, or determined as an outcome of exogenous chance events. Therefore, determining the best action for an individual in a game, in general, requires a joint analysis of every individual's decision problem.

In the previous section we have analyzed situations in which this problem could be circumvented, and hence we could analyze the problem by only considering it from the perspective of a single individual. If, independent of the other players' actions, the individual in question has an optimal action, then rationality requires taking that action, and hence we can analyze that individual's decision making problem in isolation from that of others. If every individual is in a similar situation this leads to (weakly or strictly) dominant strategy equilibrium. Remember that, the only assumptions that we used to justify dominant strategy equilibrium concept was the rationality of players (and the knowledge of own payoff function, of course). Unfortunately, many interesting games do not have a dominant strategy equilibrium and this forces us to increase the rationality requirements for individuals. The second solution concept

that we introduced, i.e., iterated elimination of dominated strategies, did just that. It required not only the rationality of each individual and the knowledge of own payoff functions, but also the (common) knowledge of other players' rationality and payoff functions. However, in this case we run into other problems: there may be too many outcomes that survive IESD actions, or different outcomes may arise as outcomes that survive IEWD actions, depending on the order of elimination.

In this section we will analyze by far the most commonly used equilibrium concept for strategic games, i.e., the Nash equilibrium concept, which overcomes some of the problems of the solution concepts introduced before.¹ As we have mentioned above, the presence of interaction among players requires each individual to form a belief regarding the possible actions of other individuals. Nash equilibrium is based on the premises that (i) each individual acts rationally given her beliefs about the other players' actions, and that (ii) these beliefs are correct. It is the second element which makes this an equilibrium concept. It is in this sense we may regard Nash equilibrium outcome as a steady state of a strategic interaction. Once every individual is acting in accordance with the Nash equilibrium, no one has an incentive to unilaterally deviate and take another action. More formally, we have the following definition:

Definition. A **Nash equilibrium** of a game G in strategic form is defined as any outcome $(a_1^*, \dots, a_n^*)$ such that

$$u_i(a_i^*, a_{-i}^*) \geq u_i(a_i, a_{-i}^*) \quad \text{for all } a_i \in A_i.$$

holds for each player i . The set of all Nash equilibria of G is denoted $\mathbf{N}(G)$.

In a two player game, for example, an action profile (a_1^*, a_2^*) is a Nash equilibrium if the following two conditions hold

$$\begin{aligned} a_1^* &\in \arg \max_{a_1 \in A_1} u_1(a_1, a_2^*) \\ a_2^* &\in \arg \max_{a_2 \in A_2} u_2(a_1^*, a_2). \end{aligned}$$

Therefore, we may say that, in a Nash equilibrium, each player's choice of action is a best response to the actions actually taken by his opponents. This suggests, and sometimes more useful, definition of Nash equilibrium, based on the notion of the best response correspon-

¹The discovery of the basic idea behind the Nash equilibrium goes back to the 1838 work of Augustin Cournot. (Cournot's work is translated into English in 1897 as *Researches into the Mathematical Principles of the Theory of Wealth*, New York: MacMillan.) The formalization and rigorous analysis of this equilibrium concept was not given until the seminal 1950 work of the mathematician John Nash. Nash was awarded the Nobel prize in economics in 1994 (along with John Harsanyi and Reinhardt Selten) for his contributions to game theory. For an exciting biography of Nash, we refer the reader to S. Nasar (1998), *A Beautiful Mind*, New York: Simon and Schuster.

dence.² We define the *best response correspondence* of player i in a strategic form game as the correspondence $B_i : A_{-i} \rightrightarrows A_i$ given by

$$\begin{aligned} B_i(a_{-i}) &= \{a_i \in A_i : u_i(a_i, a_{-i}) \geq u_i(b_i, a_{-i}) \text{ for all } b_i \in A_i\} \\ &= \arg \max_{a_i \in A_i} u_i(a_i, a_{-i}). \end{aligned}$$

(Notice that, for each $a_{-i} \in A_{-i}$, $B_i(a_{-i})$ is a set which may or may not be a singleton.) So, for example, in a 2-person game, if player 2 plays a_2 , player 1's best choice is to play some action in $B_1(a_2)$,

$$B_1(a_2) = \{a_1 \in A_1 : u_1(a_1, a_2) \geq u_1(b_1, a_2) \text{ for all } b_1 \in A_1\}.$$

For instance, in the game

	L	M	R
U	1,0	1,2	0,2
D	0,3	1,1	2,0

we have $B_1(L) = \{U\}$, $B_1(M) = \{U, D\}$ and $B_1(R) = \{D\}$, while $B_2(U) = \{M, R\}$ and $B_2(D) = \{L\}$.

The following is an easy but useful observation.

Proposition B. *For any 2-person game in strategic form G , we have $(a_1^*, a_2^*) \in \mathbf{N}(G)$ if, and only if,*

$$a_1^* \in B_1(a_2^*) \quad \text{and} \quad a_2^* \in B_2(a_1^*).$$

Exercise. Prove Proposition B.

Proposition B suggests a way of computing the Nash equilibria of strategic games. In particular, when the best response correspondence of the players are single-valued, then Proposition B tells us that all we need to do is to solve two equations in two unknowns to characterize the set of all Nash equilibria (once we have found B_1 and B_2 , that is). The following examples will illustrate.

Example. We have $\mathbf{N}(\text{BoS}) = \{(m, m), (o, o)\}$. Indeed, in this game, $B_1(o) = \{o\}$, $B_1(m) = \{m\}$, $B_2(o) = \{o\}$, and $B_2(m) = \{m\}$. These observations also show that (m, o) and (o, m) are not equilibrium points of BoS. Similar computations yield $\mathbf{N}(\text{CG}) = \{(l, l), (r, r)\}$ and $\mathbf{N}(\text{MW}) = \emptyset$.

²*Mathematical Reminder:* Recall that a function f from a set A to a set B assigns to each $x \in A$ one and only one element $f(x)$ in B . By definition, a *correspondence* f from A to B , on the other hand, assigns to each $x \in A$ a *subset* of B , and in this case we write $f : A \rightrightarrows B$. (For instance, $f : [0, 1] \rightrightarrows [0, 1]$ defined as $f(x) = \{y \in [0, 1] : x \leq y\}$ is a correspondence; draw the graph of f .) In the special case where a correspondence is single-valued (i.e. $f(x)$ is a singleton set for each $x \in A$), then f can be thought of as a function.

An easy way of finding Nash equilibrium in two-person strategic form games is to utilize the best response correspondences and the bimatrix representation. You simply have to mark the best response(s) of each player given the action choice of the other player and any action profile at which both players are best responding to each other is a Nash equilibrium. In the BoS game, for example, given player 1 plays m, the best response of player 2 is to play m, which is expressed by underscoring player 2's payoff at (m,m), and her best response to o is o, which is expressed by underscoring her payoff at (o,o).

	m	o
m	<u>2</u> , <u>1</u>	0,0
o	0,0	0, <u>1</u>

The same procedure is applied to player 1 as well. The set of Nash equilibrium is then the set of outcomes at which both players' payoffs are underscored, i.e., $\{(m,m), (o,o)\}$. $\square$

Nash equilibrium concept has been motivated in many different ways, mostly on an informal basis. We will now give a brief discussion of some of these motivations:

Self Enforcing Agreements. Let us assume that two players debate about how they should play a given 2-person game in strategic form through preplay communication. If no binding agreement is possible between the players, then what sort of an agreement would they be able to implement, if any? Clearly, the agreement (whatever it is) should be “self enforcing” in the sense that no player should have a reason to deviate from her promise if she believes that the other player will keep his end of the bargain. A Nash equilibrium is an outcome that would correspond to a self enforcing agreement in this sense. Once it is reached, no individual has an incentive to deviate from it unilaterally.

Social Conventions. Consider a strategic interaction played between two players, where player 1 is randomly picked from a population and player 2 is randomly picked from another population. For example, the situation could be a bargaining game between a randomly picked buyer and a randomly picked seller. Now imagine that this situation is repeated over time, each iteration being played between two randomly selected players. If this process settles down to an action profile, that is if time after time the action choices of players in the role of player 1 and those in the role of player 2 are always the same, then we may regard this outcome as a convention. Even if players start with arbitrary actions, as long as they remember how the actions of the previous players fared in the past and choose those actions that are better, any social convention must correspond to a Nash equilibrium. If an outcome is not a Nash equilibrium, then at least one of the players is not best responding, and sooner or later a player in that role will happen to land on a better action which will then be adopted by the players afterwards. Put differently, an outcome which is not a Nash equilibrium lacks a certain sense of *stability*, and thus if a convention were to develop about how to play a given game through time, we would expect this convention to correspond to a Nash equilibrium of the game.

Focal Points. Focal points are outcomes which are distinguished from others on the basis of some characteristics which are not included in the formalism of the model. Those characteristics may distinguish an outcome as a result of some psychological or social process and may even seem trivial, such as the names of the actions. Focal points may also arise due to the optimality of the actions, and Nash equilibrium is considered focal on this basis.

Learned Behavior. Consider two players playing the same game repeatedly. Also suppose that each player simply best responds to the action choice of the other player in the previous interaction. It is not hard to imagine that over time their play may settle on an outcome. If this happens, then it has to be a Nash equilibrium outcome. There are, however, two problems with this interpretation: (1) the play may never settle down, (2) the repeated game is different from the strategic form game that is played in each period and hence it cannot be used to justify its equilibrium.

So, whichever of the above parables one may want to entertain, they all seem to suggest that, if a reasonable outcome of a game in strategic form exists, it must possess the property of being a Nash equilibrium. In other words, being a Nash equilibrium is a *necessary* condition for being a reasonable outcome. But notice that this is a one-way statement; it would not be reasonable to claim that *any* Nash equilibrium of a given game corresponds to an outcome that is likely to be observed when the game is actually played. (More on this shortly.)

We will now introduce two other celebrated strategic form games to further illustrate the Nash equilibrium concept.

Example. Stag Hunt (SH) Two hungry hunters go to the woods with the aim of catching a stag, or at least a hare. They can catch a stag only if they both remain alert and devote their time and energy to catching it. Catching a hare is less demanding and does not require the cooperation of the other hunter. Each hunter prefers half a stag to a hare. Letting S denote the action of going after the stag, and H the action of catching a hare, we can represent this game by the following bimatrix

	S	H
S	2,2	0,1
H	1,0	1,1

One can easily verify that $\mathbf{N}(\text{SH}) = \{(S,S), (H,H)\}$.

Exercise. Hawk-Dove (HD) Two animals are fighting over a prey. The prey is worth v to each player, and the cost of fighting is c_1 for the first animal (player 1) and c_2 for the second animal (player 2). If they both act aggressively (hawkish) and get into a fight, they share the prey but suffer the cost of fighting. If both act peacefully (dovish), then they get to share the prey without incurring any cost. If one acts dovish and the other hawkish, there is no fight and the latter gets the whole prey.

(1) Write down the strategic form of this game

(2) Assume v, c_1, c_2 are all non-negative and find the Nash equilibria of this game in each of the following cases: (a) $c_1 > v/2, c_2 > v/2$, (b) $c_1 > v/2, c_2 < v/2$, (c) $c_1 < v/2, c_2 < v/2$.

We have previously introduced a simple Cournot duopoly model and analyzed its outcome by applying IESD actions. Let us now try to find its Nash equilibria. We will first find the best response correspondence of firm 1. Given that firm 2 produces $Q_2 \in [0, a/b]$, the best response of firm 1 is found by solving the first order condition

$$\frac{du_1}{dQ_1} = (a - c) - 2bQ_1 - bQ_2$$

which yields $Q_1 = \frac{a-c}{2b} - \frac{Q_2}{2}$. (Second order condition checks since $\frac{d^2u_1}{dQ_1^2} = -2b < 0$.) But notice that this equation yields $Q_1 < 0$ if $Q_2 > \frac{a-c}{b}$ while producing a negative quantity is not feasible for firm 1. Consequently, we have

$$B_1(Q_2) = \begin{cases} \frac{a-c}{2b} - \frac{Q_2}{2} & \text{if } Q_2 \leq \frac{a-c}{b} \end{cases}$$

and

$$B_1(Q_2) = \{0\} \quad \text{if } Q_2 > \frac{a-c}{b}.$$

By using symmetry, we also find

$$B_2(Q_1) = \begin{cases} \frac{a-c}{2b} - \frac{Q_1}{2}, & \text{if } Q_1 \leq \frac{a-c}{b} \\ \{0\}, & \text{if } Q_1 > \frac{a-c}{b}. \end{cases}$$

Observe next that it is impossible that either firm will choose to produce more than $\frac{a-c}{b}$ in the equilibrium (why?). Therefore, by Proposition B, to compute the Nash equilibrium all we need to do is to solve the following two equations:

$$Q_2^* = \frac{a-c}{2b} - \frac{Q_1^*}{2} \quad \text{and} \quad Q_1^* = \frac{a-c}{2b} - \frac{Q_2^*}{2}.$$

Doing this, we find that the unique Nash equilibrium of this game is

$$(Q_1^*, Q_2^*) = \left(\frac{a-c}{3b}, \frac{a-c}{3b} \right).$$

(See Figure 1.) Interestingly, this is precisely the only outcome that survives the IESD actions.

An interesting question to ask at this point is if in the Cournot model it is inefficient for these firms to produce their Nash equilibrium levels of output. The answer is yes, showing that the inefficiency of decentralized behavior may surface in more realistic settings than the scenario of the prisoners' dilemma suggests. To prove this, let us entertain the possibility that firms 1 and 2 collude (perhaps forming a cartel) and act as a monopolist with the proviso that




Figure 1: Nash Equilibrium of Cournot Duopoly Game

the profits earned in this monopoly will be distributed equally among the firms. Given the market demand, the objective function of the monopolist is

$$U(Q) = (a - c - bQ)Q$$

where $Q = Q_1 + Q_2 \in [0, 2a/b]$. By using calculus, we find that the optimal level of production for this monopoly is $Q = \frac{a-c}{2b}$. (Since the cost functions of the individual firms are identical, it does not really matter how much of this production takes place in whose plant.) Consequently,

$$\frac{\text{profits of the monopolist}}{2} = \frac{1}{2} \left(a - c - b \left(\frac{a-c}{2b} \right) \right) \left(\frac{a-c}{2b} \right) = \frac{(a-c)^2}{4b}$$

while

$$\text{profits of firm } i \text{ in the equilibrium} = u_i(Q_1^*, Q_2^*) = \frac{(a-c)^2}{9b}.$$

Thus, while both parties could be strictly better off had they formed a cartel, the equilibrium predicts that this will not take place in actuality. (Do you think this insight generalizes to the n -firm case?) $\square$

Remark. There is reason to expect that symmetric outcomes will materialize in symmetric games since in such games all agents are identical to one another. Consequently, symmetric equilibria of symmetric games is of particular interest. Formally, we define a *symmetric equilibrium* of a symmetric game as a Nash equilibrium of this game in which all players play the same action. (Note that this concept does not apply to asymmetric games.) For instance, in the Cournot duopoly game above, (Q_1^*, Q_2^*) corresponds to a symmetric equilibrium. More

generally, *if the Nash equilibrium of a symmetric game is unique, then this equilibrium must be symmetric.* Indeed, suppose that G is a symmetric 2-person game in strategic form with a unique equilibrium and $(a_1^*, a_2^*) \in \mathbf{N}(G)$. But then using the symmetry of G one may show easily that (a_2^*, a_1^*) is a Nash equilibrium of G as well. Since there is only one equilibrium of G , we must then have $a_1^* = a_2^*$. $\square$

Nash equilibrium requires that no individual has an incentive to deviate from it. In other words, it is possible that at a Nash equilibrium a player may be indifferent between her equilibrium action and some other action, given the other players' actions. If we do not allow this to happen, we arrive at the notion of a **strict Nash equilibrium**. More formally, an action profile a^* is a strict Nash equilibrium if

$$u_i(a_i^*, a_{-i}^*) > u_i(a_i, a_{-i}^*) \quad \text{for all } a_i \in A_i \text{ such that } a_i \neq a_i^*$$

holds for each player i .

For example, both Nash equilibria are strict in Stag-Hunt game, whereas the unique equilibrium of the following game, (M,R), is not strict

	L	R
T	-1, 0	0, -1
M	0, 1	0, 1
B	1, -1	-1, 0

The Nash Equilibrium and Dominant/Dominated Actions

Now that we have seen a few solution concepts for games in strategic form, we should analyze the relations between them. We turn to such an analysis in this section.

It follows readily from the definitions that every strictly dominant strategy equilibrium is a weakly dominant strategy equilibrium, and every weakly dominant strategy equilibrium is a Nash equilibrium. Thus,

$$\mathbf{D}^s(G) \subseteq \mathbf{D}^w(G) \subseteq \mathbf{N}(G)$$

for all strategic games G . For instance, (C,C) is a Nash equilibrium for PD; in fact this is the only Nash equilibrium of this game (do you agree?).

Exercise. Show that if all players have a strictly dominant strategy in a strategic game, then this game must have a unique Nash equilibrium.

However, there may exist a Nash equilibrium of a game which is not a weakly or strictly dominant strategy equilibrium; the BoS provides an example to this effect. What is more interesting is that a player may play a weakly dominated action in Nash equilibrium. Here is an example:

	α	β	
α	0,0	1,0	(1)
β	0,1	3,3	

Here (α, α) is a Nash equilibrium, but playing β weakly dominates playing α for both players. This observation can be stated in an alternative way:

Proposition C. *A Nash equilibrium need not survive the IEWD actions.*

Yet the following result shows that if IEWD actions somehow yields a unique outcome, then this must be a Nash equilibrium in finite strategic games.

Proposition D. *Let G be a game in strategic form with finite action spaces. If the iterated elimination of weakly dominated actions results in a unique outcome, then this outcome must be a Nash equilibrium of G .³*

Proof. For simplicity, we provide the proof for the 2-person case, but it is possible to generalize the argument in a straightforward way. Let the only actions that survive the IEWD actions be a_1^* and a_2^* , but to derive a contradiction, suppose that $(a_1^*, a_2^*) \notin \mathbf{N}(G)$. Then, one of the players must not be best-responding to the other, say this player is the first one. Formally, we have

$$u_1(a_1^*, a_2^*) < u_1(a_1', a_2^*) \quad \text{for some } a_1' \in A_1. \quad (2)$$

³So, for instance, $(1, \dots, 1)$ must be a Nash equilibrium of guess-the average game.

But a'_1 must have been weakly dominated by some other action $a''_1 \in A_1$ at some stage of the elimination process, so

$$u_1(a'_1, a_2) \leq u_1(a''_1, a_2) \quad \text{for each } a_2 \in A_2 \text{ not yet eliminated at that stage.}$$

Since a_2^* is never eliminated (by hypothesis), we then have

$$u_1(a'_1, a_2^*) \leq u_1(a''_1, a_2^*).$$

Now if $a''_1 = a_1^*$, then we contradict (2). Otherwise, we continue as we did after (2) to obtain an action $a'''_1 \notin \{a'_1, a''_1\}$ such that $u_1(a'_1, a_2^*) \leq u_1(a'''_1, a_2^*)$. If $a'''_1 = a_1^*$ we are done again, otherwise we continue this way and eventually reach the desired contradiction since A_1 is a finite set by hypothesis. ■

However, even if IEWD actions results in a unique outcome, there may be Nash equilibria which do not survive IEWD actions (The game given by (1) illustrates this point). Furthermore, it is important that IEWD actions leads to a unique outcome for the proposition to hold. For example in the BoS game all outcomes survive IEWD actions, yet the only Nash equilibrium outcomes are (m,m) and (o,o). One can also, by trivially modifying the proof given above show that if IESD actions results in a unique outcome, then that outcome must be a Nash equilibrium. In other words, any finite and dominance solvable game has a unique Nash equilibrium. But how about the converse of this? Is it the case that a Nash equilibrium always survives the IESD actions. In contrast to the case with IEWD actions (recall Proposition C), the answer is given in the affirmative by our next result.

Proposition E. *Let G be a 2-person game in strategic form. If $(a_1^*, a_2^*) \in \mathbf{N}(G)$, then a_1^* and a_2^* must survive the iterated elimination of strictly dominated actions.*

Proof. We again give the proof in the 2-person case for simplicity. To obtain a contradiction, suppose that $(a_1^*, a_2^*) \in \mathbf{N}(G)$, but either a_1^* or a_2^* is eliminated at some iteration. Without loss of generality, assume that a_1^* is eliminated before a_2^* . Then, there must exist an action $a'_1 \in A_1$ (not yet eliminated at the iteration at which a_1^* is eliminated) such that,

$$u_1(a_1^*, a_2) < u_1(a'_1, a_2) \quad \text{for each } a_2 \in A_2 \text{ not yet eliminated.}$$

But a_2^* is not yet eliminated, and thus

$$u_1(a_1^*, a_2^*) < u_1(a'_1, a_2^*)$$

so that (a_1^*, a_2^*) cannot be a Nash equilibrium, a contradiction. ■

Difficulties with the Nash Equilibrium

Given that the Nash equilibrium is the most widely used equilibrium concept in economic applications, it is important to understand its limitations. We discuss some of these as the final order of business in this chapter.

(1) A Nash equilibrium may involve a weakly dominated action by some players.

We observed this possibility in Proposition C. Ask yourself if (α, α) in the game (1) is a sensible outcome at all. You may say that if player 1 is “certain” that player 2 will play α and vice versa, then it is. But if either one of the players assigns a probability in her mind that her opponent may play β , the expected utility maximizing (rational) action would be to play β , *no matter how small this probability is*. Since it is rare that all players are “certain” about the intended plays of their opponents (even if pre-play negotiation is possible), weakly dominated Nash equilibrium appears unreasonable. This leads us to *refine* the Nash equilibrium in the following manner.

Definition. An **undominated Nash equilibrium** of a game G in strategic form is defined as any Nash equilibrium $(a_1^*, \dots, a_n^*)$ such that none of the a_i^* s is a weakly dominated action. The set of all undominated Nash equilibria of G is denoted $\mathbf{N}_{\text{undom}}(G)$.

Example. If G denotes the game given in (1), then $\mathbf{N}_{\text{undom}}(G) = \{(\beta, \beta)\}$. On the other hand, $\mathbf{N}_{\text{undom}}(G) = \mathbf{N}(G)$ where $G = \text{PD}, \text{BoS}, \text{CG}$. The same equality holds for the linear Cournot model. (*Question:* Are all strict Nash equilibria of a game in strategic form undominated?)

Exercise. Compute the set of all Nash and undominated Nash equilibria of the chairman’s paradox game.

(2) Nash equilibrium need not exist.

For instance, $\mathbf{N}(\text{MW}) = \emptyset$. Thus the notion of Nash equilibrium does not help us predict how the MW game would be played in practice. However, it is possible to circumvent this problem to some extent by enlarging the set of actions available to the players by allowing them to “randomize” among their actions. This leads us to the notion of a *mixed strategy* which we shall talk about later in the course.

(3) Nash equilibrium need not be unique.

The BoS and CG provide two examples to this effect. This is a troubling issue in that multiplicity of equilibria avoids making a sharp prediction with regard to the actual play of the game. (What do you think will be the outcome of BoS?) However, sometimes preplay negotiation and/or conventions may provide a way out of this problem.

Preplay Negotiation. Consider the CG game and allow the players to communicate (cheap talk) prior to the game being played. What do you think will be the outcome then? Most people answer this question as (r,r). The reason is that agreement on the outcome (r,r) seems in the nature of things, and what is more, there is no reason why players should not play r once this agreement is reached (i.e. such an agreement is *self-enforcing*). Thus, pure coordination games like CG can often be “solved” via preplay negotiation. (More on this shortly.)

But how about BoS? It is not at all obvious which agreement would surface in the preplay communication in this game, and hence, even if an agreement on either (m,m) or (o,o) would be self-enforcing, preplay negotiation does not help us “solve” the BoS. Maybe we should learn to live with the fact that some games do not admit a natural “solution.”

Focal Points. It has been argued by many game theorists that the story of some games isolate certain Nash equilibria as “focal” in that certain details that are not captured by the formalism of a game in strategic form may actually entail a clear path of play. The following will illustrate.

Example. (*A Nash Demand Game*) Suppose that two individuals (1 and 2) face the problem of dividing \$100 among themselves. They decide to use the following method in doing this: each of them will simultaneously declare how much of the \$100 (s)he wishes to have, and if their total demand exceeds \$100 no one will get anything (the money will then go to a charity) while they will receive their demands otherwise (anything left on the table will go to a charity).

We may formulate this scenario as a 2-person game in strategic form where $A_i = [0, 100]$ and

$$u_i(x_1, x_2) = \begin{cases} x_i, & \text{if } x_1 + x_2 \leq 100 \\ 0, & \text{otherwise.} \end{cases}$$

Notice that we are assuming here that money is utility; an assumption which is often useful. (*Caveat:* But this is not an unexceptionable assumption - what if the bargaining was between a father and his 5 year old daughter or between two individuals who hate each other?).

- Play the game.
- Verify that the set of Nash equilibria of this game is

$$\{(x_1, x_2) \in [0, 100]^2 : x_1 + x_2 = 100\}.$$

- Well, there are just too many equilibria here; any division of \$100 is an equilibrium! Thus, for this game, the predictions made on the basis of the Nash equilibrium are bound to be very weak. Yet, when people actually played this game in the experiments, in an overwhelming number of times the outcome (50, 50) is observed to surface. So, in

this example, 50-50 split appears to be a focal point suggesting that *equity* considerations (which are totally missed by the formalism of the game theory we have developed so far) may play a role in certain Nash equilibrium to be selected in actual play. $\square$

Unfortunately, the notion of a focal point is an elusive one. It is difficult to come up with a theory for it since it is not clear what is the general principle that underlies it. The above example provides, after all, only a single instance of it; one can think of other scenarios with a focal equilibrium.⁴ It is our hope that experimental game theory (which we shall talk about further later on) will shed light into the matter in the future.

(4) Nash equilibrium is not immune to coalitional deviations.

Consider again the CG game in which we argued that preplay negotiation would eliminate the Nash equilibrium (1,1). The idea is that the players can *jointly* deviate from the outcome (1,1) through communication that takes place prior to play), for at the Nash equilibrium outcome (r,r) they are both strictly better off. This suggests the following refinement of the Nash equilibrium.

Definition. A **Pareto optimal Nash equilibrium** of a game G in strategic form is any Nash equilibrium $a^* = (a_1^*, \dots, a_n^*)$ such that there does not exist another equilibrium $b^* = (b_1^*, \dots, b_n^*) \in N(G)$ with

$$u_i(a^*) < u_i(b^*) \quad \text{for each } i \in N.$$

We denote the set of all Pareto optimal Nash equilibrium of G by $N_{PO}(G)$.

A Pareto optimal Nash equilibrium outcome in a 2-person game in strategic form is particularly appealing (when preplay communication is allowed), for once such an outcome has been somehow realized, the players would not have an incentive from deviating from it neither unilaterally (as the Nash property requires) nor jointly (as Pareto optimality requires). As you would expect, this refinement of Nash equilibrium delivers us what we wish to find in the CG: $N_{PO}(CG) = \{(r,r)\}$. As you might expect, however, the Pareto optimal Nash equilibrium concept does not help us “solve” the BoS, for we have $N_{PO}(BoS) = N(BoS)$.

⁴Here is another game in strategic form with some sort of a focal point. Two players are supposed to partition the letters A,B,C,D,E,F,G,H with the proviso that player 1’s list must contain A and player 2’s list must contain H. If their lists do not overlap, then they both win, they lose otherwise. (How would you play this game in the place of player 1? Player 2?) What happens very often when the game is played in the experiments is that people in the position of player 1 chooses {A,B,C,D} and people in the position of player 2 chooses {E,F,G,H}; what is going on here, how do people coordinate so well? For more examples of this sort and a thorough discussion of focal points, an excellent reference is T. Schelling (1960), *The Strategy of Conflict*, London: Oxford University Press.

The fact that Pareto optimal Nash equilibrium refines the Nash equilibrium points to the fact that the latter is not immune to *coalitional deviations*. This is because the stability achieved by the Nash equilibrium is by means of avoiding only the unilateral deviations of each individual. Put differently, the Nash equilibrium does not ensure that no coalition of the players will find it beneficial to defect. The Pareto optimal Nash equilibrium somewhat corrects for this through avoiding defection of the entire group of the players (the so-called grand coalition) in addition to that of the individuals (the singleton coalitions). Unfortunately, this refinement does not solve the problem entirely. Here is a game in which the Pareto optimal Nash equilibrium does *not* refine the Nash equilibrium in a way that deals with coalitional considerations in a satisfactory way.

Example. In the following game G player 1 chooses rows, player 2 chooses columns and player 3 chooses tables.

	α	β	
a	1,1,-5	-5,-5,0	if player 3 chooses U
b	-5,-5,0	0,2,7	

	α	β	
a	1,1,6	-5,-5,0	if player 3 chooses D
b	-5,-5,0	-2,-2,0	

(For instance, we have $N = \{1, 2, 3\}$, $A_3 = \{U, D\}$ and $u_3(a, \beta, D) = 0$.) In this game we have

$$\mathbf{N}_{PO}(G) = \{(b, \beta, U), (a, \alpha, D)\} = \mathbf{N}(G),$$

but coalitional considerations indicate that the equilibrium (a, α, D) is rather unstable, provided that players can communicate prior to play. Indeed, it is quite conceivable in this case that players 2 and 3 would form a coalition and deviate from (a, α, D) equilibrium by publicly agreeing to take actions β and U, respectively. Since this is clearly a self-enforcing agreement, it casts doubt on the claim that (a, α, D) is a reasonable prediction for this game. $\square$

You probably see where the above example is leading to. It suggest that there is merit in refining even the Pareto optimal Nash equilibrium by isolating those Nash equilibria that are immune against *all* possible coalitional deviations. To introduce this idea formally, we need a final bit of

Notation. Let $A = \times_{i \in N} A_i$ be the outcome space of an n -person game in strategic form, and let $(a_1, \dots, a_n) \in A$. For each $K \subseteq N$, we let a_K denote the vector $(a_i)_{i \in K} \in \times_{i \in K} A_i$, and a_{-K} the vector $(a_i)_{i \in N \setminus K} \in \times_{i \in N \setminus K} A_i$. By (a_K, a_{-K}) , we then mean the outcome $(a_1, \dots, a_n)$. Clearly, a_K is the profile of actions taken by all players who belong to the coalition K , and we denote the set of all such profiles by A_K (that is, $A_K = \times_{i \in K} A_i$ by definition). Similarly,

a_{-K} is the profile of actions taken by all players who does not belong to K , and A_{-K} is a shorthand notation for the set $A_{-K} = \times_{i \in N \setminus K} A_i$.

Definition. A **Strong Nash equilibrium** of a game G in strategic form is any outcome $a^* = (a_1^*, \dots, a_n^*)$ such that, for all nonempty coalitions $K \subseteq N$ and all $a_K \in A_K$, there exists a player $i \in K$ such that

$$u_i(a_K^*, a_{-K}^*) \geq u_i(a_K, a_{-K}^*).$$

We denote the set of all strong Nash equilibrium of G by $\mathbf{N}_S(G)$.⁵

While its formal definition is a bit mouthful, all that the strong Nash equilibrium concept does is to choose those outcomes at which no coalition can find it in the interest of *each* of its members to deviate. Clearly, we have

$$\mathbf{N}_S(G) \subseteq \mathbf{N}_{PO}(G) \subseteq \mathbf{N}(G)$$

for any game G in strategic form. Since, for 2-person games the notions of Pareto optimal and strong Nash equilibrium coincide (why?), the only strong Nash equilibrium of the CG is (r,r). On the other hand, in the 3-person game discussed above, we have $\mathbf{N}_S(G) = \{(b, \beta, U)\}$ as is desired (verify!).

Unfortunately, while the notion of the strong Nash equilibrium solves some of our problems, it is itself not free of difficulties. In particular, in many interesting games no strong Nash equilibrium exists, for it is simply too demanding to disallow for *all* coalitional deviations. What we need instead is a theory of coalition formation so that we can look for the Nash equilibria that are immune to deviations by those coalitions that are likely to form. At present, however, there does not exist such a theory that is commonly used in game theory, the issue awaits much further research.⁶

⁵The notion of the strong Nash equilibrium was first introduced by the mathematician and economist Robert Aumann.

⁶If you are interested in coalitional refinements of the Nash equilibrium, a good place to start is the highly readable paper by D. Bernheim, B. Peleg and M. Whinston (1987), "Coalition-proof Nash equilibria I: Concepts," *Journal of Economic Theory*, 42, pp. 1-12.

Nash Equilibrium: Applications

Prof. Levent Koçkesen

Columbia University

and

Prof. Efe A. Ok

New York University

Introduction

In this section, we shall consider several economic scenarios which are modeled well by means of strategic games. We shall also examine the predictions that game theory provides in such scenarios by using some of the equilibrium concepts that we have studied so far. One major objective of this section is actually to establish a solid understanding of the notion of Nash equilibrium, undoubtedly the most commonly used equilibrium concept in game theory. We contend that the best way of understanding the pros and cons of Nash equilibrium is seeing this concept in action. For this reason we shall consider below quite a number of examples. Most of these examples are the toy versions of more general economic models and we shall return to some of them in later chapters when we are better equipped to cover more realistic scenarios.

Auctions

Many economic transactions are conducted through auctions. Governments sell treasury bills, foreign exchange, mineral rights, and more recently airwave spectrum rights via auctions. Art work, antiques, cars, and houses are also sold by auctions. Auction theory has also been applied to areas as diverse as queues, wars of attrition, and lobbying contests.¹

There are four commonly used and studied forms of auctions: the ascending-bid auction (also called English auction), the descending-bid auction (also called Dutch auction), the first-price sealed bid auction, and the second-price sealed bid auction (also known as Vickrey auction²). In the *ascending-bid auction*, the price is raised until only one bidder remains, and that bidder wins the object at the final price. In the *descending-bid auction*, the auctioneer starts at a very high price and lowers it continuously until the someone accepts the currently announced price. That bidder wins the object at that price. In the *first-price sealed bid auction* each bidder submits her bid in a sealed envelope without seeing others' bids, and the object is sold to the highest bidder at her bid. The *second-price sealed bid auction* works the same way except that the winner pays the second highest bid.

¹For a good introductory survey to the auction theory see Paul Klemperer (1999), "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys*, 13(3), July 1999, pp. 227-286.

²Named after William Vickrey of Columbia University who was awarded the Nobel Prize in economics in 1996.

In this section we will analyze the last two forms of auctions, not only because they are simpler to analyze but also because under the assumptions we will work with in this section the first-price sealed bid auction is strategically equivalent to descending bid auction and the second-price sealed bid auction is strategically equivalent to ascending bid auction.

For simplicity we will assume there are only two individuals, players 1 and 2, who are competing in an auction for a valuable object. While this may require a stretch of imagination, it is commonly known that the value of the object to the player i is v_i dollars, $i = 1, 2$, where $v_1 > v_2 > 0$. (What we mean by this is that player i is indifferent between buying the object at price v_i and not buying it.) The outcome of the auction, of course, depends on the rules of the auctioning procedure. In fact, identifying the precise nature of the outcomes in a setting like this (and in similar scenarios) under various procedures is the subject matter of a very likely subfield of game theory, namely the auction theory. In this section, our aim is to provide an elementary introduction to this topic. let us then begin with analyzing this game theoretic scenario first under the most common auctioning procedure.

First-price sealed bid auction

The rules of the first-price auction is such that after both players cast their bid (without observing each others' bid), the highest bidder wins the object and pays her own bid. In case of a tie, the object is awarded to player 1.³

Assuming that utility is money (i.e., individuals are risk neutral), this bargaining procedure results in the 2-person game in strategic form $G = (A_1, A_2, u_1, u_2)$ where $A_1 = A_2 = \mathbf{R}_+$,

$$u_1(b_1, b_2) = \begin{cases} v_1 - b_1, & \text{if } b_1 \geq b_2 \\ 0, & \text{otherwise} \end{cases}$$

and

$$u_2(b_1, b_2) = \begin{cases} v_2 - b_2, & \text{if } b_2 > b_1 \\ 0, & \text{otherwise} \end{cases}$$

for all $(b_1, b_2) \in \mathbf{R}_+^2$. (Here b_i stands for the bid of player i , $i = 1, 2$).

We now wish to identify the set of Nash equilibria of G . (In case you are wondering why we are not checking for dominant strategy equilibrium, note that the following analysis will demonstrate that $\mathbf{D}^s(G) = \mathbf{D}^w(G) = \emptyset$.) Rather than computing the best response correspondences of the players, we adopt here instead a direct approach towards this goal. Let us try to find what properties a Nash equilibrium has to satisfy. We first claim that

(1) In any Nash equilibrium player 1 (the individual who values the object the most) wins the object.

³There are other tie-breaking methods such as, randomly selecting a winner (by means of coin toss, say). Our choice of the tie-breaking rule is useful in that it leads to a simple analysis. The reader should not find it difficult to modify the results reported below by using other tie-breaking rules.

Proof: Let (b_1^*, b_2^*) be a Nash equilibrium, but for a contradiction, suppose player 1 does not win the object. This implies that $b_1^* < b_2^*$ and player 1's payoff in equilibrium is zero, i.e. $u_1(b_1^*, b_2^*) = 0$. Now if $b_2^* \leq v_2$, then $b_2^* < v_1$ (since $v_2 < v_1$), and hence bidding, say b_2^* , is a strictly better response for player 1 when player 2 is bidding b_2^* . Therefore, bidding a strictly smaller amount than b_2^* cannot be a best response for player 1. If, on the other hand, $b_2^* > v_2$, then $u_2(b_1^*, b_2^*) < 0$ so that bidding anything in the interval $[0, b_1^*]$ is a profitable deviation for player 2. In either case, then, we obtain a contradiction to the hypothesis that (b_1^*, b_2^*) is an equilibrium. Therefore, we conclude that in any equilibrium (b_1^*, b_2^*) of G player 1 obtains the object, that is, $b_1^* \geq b_2^*$.

Secondly,

(2) $b_1^* > b_2^*$ cannot hold in equilibrium, for in this case player 1 would deviate by bidding, say, b_2^* and increase her payoff from $v_1 - b_1^*$ to $v_1 - b_2^*$. Together with our finding that $b_1^* \geq b_2^*$, this implies that $b_1^* = b_2^*$ must hold in equilibrium.

Thirdly,

(3) Neither $b_1^* < v_2$ nor $b_1^* > v_1$ can hold (player 2 would have a profitable deviation in the first case, and player 1 in the second case).

So, any Nash equilibrium (b_1^*, b_2^*) of this game must satisfy

$$v_2 \leq b_1^* = b_2^* \leq v_1.$$

Is any pair (b_1^*, b_2^*) that satisfy these inequalities an equilibrium? Yes. The inequality $v_2 \leq b_1^*$ guarantees that player 2 does not wish to win the object when player 1 bids b_1 , so his action is optimal. The inequality, $b_2^* \leq v_1$, on the other hand, guarantees that player 1 is also best responding. We thus conclude that

$$\mathbf{N}(G) = \{(b_1, b_2) : v_2 \leq b_1 = b_2 \leq v_1\}$$

Exercise. Verify the above conclusion by means of computing the best response correspondences of the players 1 and 2, and plotting their graph in the (b_1, b_2) space.

While $\mathbf{N}(G)$ is rather a large set and hence does not lead us to a sharp prediction, refining this set by eliminating the weakly dominated actions solves this problem. Indeed, it is easily verified that bidding anything strictly higher than v_2 is a weakly dominated action for player 2. To see this, suppose player 2 bids b'_2 which is strictly higher than v_2 . Now, if player 1's bid b_1 is greater than equal to b'_2 , then player 1 wins the object and player 2's payoff to b'_2 and to bidding her valuation v_2 are both zero. If, however, player 1's bid is strictly smaller than b'_2 but greater than or equal to v_2 , then player 2 wins by bidding b'_2 but obtains a negative payoff since she pays more than her valuation. The payoff to bidding v_2 , on the other hand, is zero. Similarly, bidding v_2 is strictly better than bidding b'_2 if player 1's bid is strictly smaller than

v_2 . The following table summarizes this discussion. (Does bidding v_2 weakly dominate bids less than v_2 as well?).

	$b_1 \geq b'_2$	$v_2 \leq b_1 < b'_2$	$b_1 < v_2$
v_2	0	0	0
b'_2	0	$v_2 - b'_2 < 0$	$v_2 - b'_2 < 0$

Consequently, we have

$$\mathbf{N}_{\text{undom}}(G) = \{(v_2, v_2)\}.$$

Now, there is an intriguing normative problem with this equilibrium: the first player is not bidding his true valuation. It is often argued that it would be desirable to design an auctioning method in which all players are induced to bid their true valuations in equilibrium. But is such a thing possible? This question was answered in the affirmative by the economist William Vickrey who has showed that truth-telling can be established even as a dominant action by modifying the rules of the auction suitably. Let us carefully examine Vickrey's modification.

Second-price sealed bid (Vickrey) auction

The rules of the second-price auction is such that after both players cast their bid (without observing each others' bid), the highest bidder wins the object and pays the bid of the other player. In case of a tie, the object is awarded to player 1.

Assuming that utility is money, this bargaining procedure results in the 2-person game in strategic form $G' = (A_1, A_2, u_1, u_2)$ where $A_1 = A_2 = \mathbf{R}_+$,

$$u_1(b_1, b_2) = \begin{cases} v_1 - b_2, & \text{if } b_1 \geq b_2 \\ 0, & \text{otherwise} \end{cases}$$

and

$$u_2(b_1, b_2) = \begin{cases} v_2 - b_1, & \text{if } b_2 > b_1 \\ 0, & \text{otherwise} \end{cases}$$

for all $(b_1, b_2) \in \mathbf{R}_+^2$. (Contrast G' with the game G we studied above.)

We now claim that $\mathbf{D}^w(G') = \{(v_1, v_2)\}$. To see that bidding $b_1 = v_1$ is a dominant action for player 1, we distinguish between two cases:

Case 1. Player 2 bids strictly less than v_1 (that is, $b_2 < v_1$)

In this case, by bidding v_1 player 1 wins the object and achieves a utility level of $v_1 - b_2 > 0$. Bidding strictly less than v_1 either makes her win the object (if $b_2 \leq b_1 < v_1$) with payoff $v_1 - b_2$, or she loses the object (if $b_1 < b_2 < v_1$) with payoff zero. So, bidding v_1 is at least as good as bidding strictly less than v_1 , and sometimes it is strictly better. Bidding strictly greater than v_1 , on the other hand, brings player 1 a payoff of $v_1 - b_2 > 0$, the same payoff as she would get by bidding v_1 .

Case 2. Player 2 bids v_1 (that is, $b_2 = v_1$)

In this case, every bid brings player 1 a payoff of zero.

Case 3. Player 2 bids strictly more than v_1 (that is, $b_2 > v_1$)

In this case, player 1 loses the object and obtains utility 0. So bidding v_1 is again optimal for player 1 since winning the object in this case would entail negative utility for him.

Consequently, bidding v_1 is a dominant action for player 1. A similar reasoning shows that bidding v_2 is a dominant action for player 2, and hence we have $\mathbf{D}^w(G') = \{(v_1, v_2)\}$ as is sought. We hope you agree that this is a very nice result. Since, a weakly dominant strategy equilibrium is also a Nash equilibrium, we also have that (v_1, v_2) is a Nash equilibrium. However, there are other Nash equilibria of this game. For example $(v_1, 0)$ is a Nash equilibrium too (verify).

Exercise. Generalize the above analysis by considering $n \geq 2$ many individuals assuming that the value of the object to player i is v_i dollars, $i = 1, \dots, n$, where $v_1 > \dots > v_n$, that the object is given to the highest bidder with the smallest index in both the first and second-price auctions, and that the winner pays the second highest bid in the second-price auction.

By an ingenious modification of the first-price auction, therefore, Vickrey was able to guarantee the truthful revelation of the preferences of the agents in dominant strategy equilibrium. This result shows that, by designing the rules of interaction carefully, one may force the individuals to coordinate on normatively appealing outcomes, and this even without knowing the true valuations of the individuals! Vickrey's technique provides a foundation for the theory of implementation which has important applications in public economics where one frequently needs to put on the mask of a *social engineering*. We shall talk about some of these applications later on.

Buyer-Seller Games

A seller, call him player s , is in possession of an object that is worth v_s dollars to him (that is, player s is indifferent between receiving v_s dollars for the object and keeping the object). The value of this object is v_b dollars to a potential buyer, player b . We assume in what follows that

$$v_b > v_s > 0.$$

So, since the value of the object is higher for the buyer than it is for the seller, an efficient state of affairs demand that trade takes place. But what is the price that player b will pay to player s ? The buyer wants to pay only v_s (well, she wants to pay nothing, but she knows that the seller will not sell the object at a price strictly less than v_s) while the seller wants to charge v_b . The actual price of the object will thus be determined through the bargaining of the players. Different bargaining scenarios would presumably lead to different equilibrium prices. To demonstrate this, we shall consider here two such scenarios.⁴

Bargaining Scenario 1: Sealed-bid first-price auction

Each party proposes a price between v_s and v_b simultaneously (by means of a sealed bid). If the price suggested by the buyer p_b is *strictly higher* than that proposed by the seller, say p_s , then trade takes place at price p_b , otherwise there is no trade. Assuming that utility is money, this bargaining procedure results in the 2-person game in strategic form (A_b, A_s, u_b, u_s) where $A_b = A_s = [v_s, v_b]$,

$$u_b(p_b, p_s) = \begin{cases} v_b - p_b, & \text{if } p_b > p_s \\ 0, & \text{otherwise} \end{cases}$$

and

$$u_s(p_b, p_s) = \begin{cases} p_b - v_s, & \text{if } p_b > p_s \\ 0, & \text{otherwise} \end{cases}$$

for all $(p_b, p_s) \in [v_s, v_b]^2$.

We first observe that there is no Nash equilibrium (p_b^*, p_s^*) in which the buyer b buys the object. Indeed, if $p_b^* > p_s^*$, then bidding anything between p_s^* and p_b^* (e.g. $p_b^*/2 + p_s^*/2$) would be a strictly better response for player b (than playing p_b^*) against p_s^* .

Exercise. Consider the game described above. Show that the best response correspondence of players b and s are given as

$$B_b(p_s) = \begin{cases} \emptyset, & \text{if } p_s \in [v_s, v_b) \\ [v_s, v_b], & \text{if } p_s = v_b \end{cases}$$

⁴It is very likely that these scenarios will strike you as unrealistic. The objective of these examples is, however, not achieving a satisfactory level of realism, but rather to illustrate the use of Nash equilibrium in certain simple buyer-seller games. In later chapters, we will return to this setting and consider much more realistic bargaining scenarios that involve sequential offers and counteroffers by the players.

and

$$B_s(p_b) = \begin{cases} [v_s, p_b], & \text{if } p_b \in (v_s, v_b] \\ [v_s, v_b], & \text{if } p_b = v_s \end{cases}$$

respectively. Deduce from this that the only equilibrium of the game is $(p_b^*, p_s^*) = (v_s, v_b)$.

Bargaining Scenario 2: Modified sealed-bid first-price auction

Each party proposes a price between v_s and v_b simultaneously (by means of a sealed bid). If the price suggested by the buyer p_b is *at least as large as* that proposed by the seller, say p_s , then trade takes place at price p_b , otherwise there is no trade. Assuming again that utility is money, this bargaining procedure results in the 2-person game in strategic form (A_b, A_s, u_b, u_s) where $A_b = A_s = [v_s, v_b]$,

$$u_b(p_b, p_s) = \begin{cases} v_b - p_b, & \text{if } p_b \geq p_s \\ 0, & \text{otherwise} \end{cases}$$

and

$$u_s(p_b, p_s) = \begin{cases} p_b - v_s, & \text{if } p_b \geq p_s \\ 0, & \text{otherwise} \end{cases}$$

for all $(p_b, p_s) \in [v_s, v_b]^2$.

If you have solved the exercise above, you will find it easy to show that we have

$$B_b(p_s) = \begin{cases} \{p_s\}, & \text{if } p_s \in [v_s, v_b) \\ [v_s, v_b], & \text{if } p_s = v_b \end{cases}$$

and

$$B_s(p_b) = \begin{cases} [v_s, p_b], & \text{if } p_b \in (v_s, v_b] \\ [v_s, v_b], & \text{if } p_b = v_s \end{cases}$$

in this game. Consequently, $(p_b^*, p_s^*) \in B_b(p_s^*) \times B_s(p_b^*)$ holds if, and only if, either $(p_b^*, p_s^*) = (v_s, v_b)$ (the no-trade equilibrium) or $p_b^* = p_s^* \in [v_s, v_b]$ (see Figure 2).

Therefore, with a minor modification of the bargaining procedure, one is able to generate many equilibria in which trade occurs. (This is an important observation for especially the seller, for, in many instances, it is the seller who design the bargaining procedure.) However, the prediction of the Nash equilibrium in the resulting game is less than satisfactory due to the large multiplicity of equilibria. (Check if undominated and/or Pareto optimal Nash equilibria provide sharper predictions here.)

Price Competition Models

Game theory has many applications in the field of industrial organization. We have already encountered one such application when we have considered in some detail the model of Cournot duopoly in the previous chapter. Recall that in this scenario individual firms were modeled as competing in the market by choosing their output levels. However, it has been argued in the literature that this model is not entirely satisfactory, especially if one is interested in the short-run decision making of the firms. For, as the argument goes, in the short-run firms would find it too costly to adjust their output level at will; it is rather through price setting that they engage in competition with other firms.⁵ To deal with this problem, several oligopoly models in which firms choose their prices (as opposed to quantities) were developed in the literature. We now briefly discuss such a price competition model, which leads to a dramatically different conclusion than does the Cournot model..

Bertrand Duopoly with Homogeneous Products

Consider the market structure underlying the linear Cournot model, but this time assume that the firms in the market engage in price competition, that is, they choose how much to charge for their products. Recalling that a is the maximum price level in the market, we thus model the action space for firm $i = 1, 2$ as $[0, a]$. The profit function of firm i on $[0, a]^2$ in this model (called the *linear Bertrand duopoly*) is defined as

$$u_i(P_1, P_2) = P_i Q_i(P_1, P_2) - c Q_i(P_1, P_2),$$

where $Q_i(P_1, P_2)$ denotes the output sold by firm i at the price profile (P_1, P_2) . If we assume that there is no qualitative difference between the products of the two firms, it would be natural to assume that the consumers always buy the cheaper good. In case both firms charge the same price, we assume that firms 1 and 2 share the market equally. These assumptions entail that

$$Q_i(P_1, P_2) = \begin{cases} \frac{a}{b} - \frac{P_i}{b}, & P_i < P_j \\ \frac{1}{2} \left(\frac{a}{b} - \frac{P_i}{b} \right), & P_i = P_j \\ 0, & P_i > P_j \end{cases}$$

where $j \neq i = 1, 2$, and complete the formulation of the model at hand as a 2-person game in strategic form.

An immediate question to ask is if our prediction (based on the Nash equilibrium) about the market outcome would be different in this model than in the linear Cournot duopoly model. The answer is easily seen to be yes. To see this, recall that in the Nash equilibrium of the linear Cournot duopoly model both firms charge the same price, namely

$$P_1 = P_2 = a - b \left(\frac{a - c}{3b} \right) = \frac{2a}{3} + \frac{c}{3},$$

⁵The argument was first given by the French mathematician Joseph Bertrand in 1883 as a critique of the Cournot model.

so that firm 1's level of profits is found as

$$u_i(P_1, P_2) = (P_1 - c)Q_1(P_1, P_2) = \left(\frac{2(a-c)}{3}\right) \frac{1}{2} \left(\frac{a}{b} - \frac{1}{b} \left(\frac{2a}{3} + \frac{c}{3}\right)\right) = \frac{1}{9b} (a-c)^2.$$

But, given that $P_2 = \frac{2a}{3} + \frac{c}{3}$, if firm 1 *undercuts* firm 2 by charging a marginally smaller price than $\frac{2a}{3} + \frac{c}{3}$, say $\frac{2a}{3} + \frac{c}{3} - \varepsilon$ where ε is some small positive number, then the profit level of firm 1 increases since this firm then grabs the entire market. Indeed, it can easily be checked that

$$\lim_{\varepsilon \searrow 0} u_i \left(\frac{2a}{3} + \frac{c}{3} - \varepsilon, \frac{2a}{3} + \frac{c}{3} \right) = \frac{2}{9b} (a-c)^2 > \frac{1}{9b} (a-c)^2 = u_i \left(\frac{2a}{3} + \frac{c}{3}, \frac{2a}{3} + \frac{c}{3} \right).$$

Thus, the Cournot prices cannot constitute an equilibrium for the linear Bertrand model. (How did we conclude this, really?) The problem is that the tie-breaking rule of the Bertrand duopoly introduces a *discontinuity* to the model allowing firms to achieve relatively large gains through small alterations of their actions.⁶

What then is the equilibrium? The analysis outlined in the previous paragraph actually brings us quite close to answering this question. First observe that neither firm would ever charge a price below c as this would yield negative profits (which can always be avoided by charging exactly c dollars for the unit product). Thus, if the price profile (P_1^*, P_2^*) is a Nash equilibrium, we must have $P_1^*, P_2^* \geq c$. Is $P_1^* > P_2^* > c$ possible? No, for in this case firm 1 would be making zero profits, and thus it would better for it to charge, say, P_2^* which will ensure positive profits given that firm 2's price is P_2^* . How about $P_1^* = P_2^* > c$? This is also impossible, because in this case either firm can unilaterally increase its profits by undercutting the other firm (just as in the discussion above) contradicting that (P_1^*, P_2^*) is a Nash equilibrium. By symmetry, $P_2^* \geq P_1^* > c$ is also impossible, and hence we conclude that at least one firm must be charging precisely its unit cost c in the equilibrium. Can we have $P_1^* > P_2^* = c$ then? No, for in this case firm 2 would not be best responding; it can increase its profits by charging, say, $P_1^*/2 + c/2$. Similarly, $P_2^* > P_1^* = c$ is not possible. The only candidate for equilibrium is thus $(P_1^*, P_2^*) = (c, c)$, and this is indeed an equilibrium as you can easily check: in the Nash equilibrium of the linear Bertrand duopoly, all firms price their products at unit cost.

This is a surprising result since it envisages that all firms operate with zero profits in the equilibrium. In fact, the equilibrium outcome here is nothing but the competitive equilibrium outcome which is justified in microeconomics only by hypothesizing a very large number of firms who act as price takers (not setters) in the industry. Here, however, we predict precisely the same outcome in equilibrium with only two price setting firms!

⁶Notice that this is the third time we are observing that the tie-breaking rule is playing an important role with regard to the nature of equilibrium. This is quite typical in many interesting strategic games, and hence, it is always a good idea to inquire into the suitability of a specific tie-breaking rule in such models.

Remark. The major culprit behind the above finding is the fundamental discontinuity that the Bertrand game possesses. Indeed, as noted earlier, it is possible in this game to alter one's action marginally (infinitesimally) and increase the associated profits significantly, given the other's action. Such games are called *discontinuous* games, and often do not possess a Nash equilibrium. For example, if we modify the linear Bertrand model so that the unit cost of firm 1, call it c_1 , exceeds that of firm 2, we obtain an asymmetric Bertrand game that does not have an equilibrium. (Exercise: Prove this.) But this is not a severe difficulty. It arises only because we take the prices as continuous variables in the classic Bertrand model. If the medium of exchange was discrete but small, then there would exist an equilibrium of this game such that the high cost firm 1 charges its unit cost (and thus make zero profits) while the low cost firm 2 would grab the entire market by charging the lowest possible price strictly below c_1 . (*Challenge:* Formalize and prove this claim.) $\square$

Spatial Voting Games

The example of Chairman's paradox was our first excursion into voting theory which provides an extensive realm for fruitful applications of game theory. In this section we provide a more daring excursion, and introduce the so-called spatial voting model.

If we think of the policy space as one-dimensional, then we can identify the set of all political positions with the closed interval $[0, 1]$. Here we may think of 0 as the most leftist position and 1 as the most rightist position. The interpretation of any other point in $[0, 1]$ is then given in the straightforward way.⁷ Let us assume next that each voter has an ideal position in the political spectrum, and evaluate every other policy in $[0, 1]$ by looking at the distance between this point and her ideal point. For instance, the voter with the ideal point $1/2 \in [0, 1]$ likes the point $1/4$ better than 1. More generally, an individual with ideal point x in $[0, 1]$ likes the point y better than z iff $|x - y| < |x - z|$. Such preferences are called *single-peaked* in the literature because, for any $x \in [0, 1]$, the mapping $y \mapsto |x - y|$ is strictly increasing on $[0, x]$ and strictly decreasing on $[x, 1]$. (Plot the graph of this mapping on $[0, 1]$ and see for yourself.)

We model the society as a continuum, and posit that the voters (which can be identified with their ideal positions in this model) are distributed uniformly over $[0, 1]$. Thus $1/2$ corresponds to the *median* ideal position in the society, that is, the ideal positions of exactly half of the society lie to the left of $1/2$.

The players in a voting game are the political candidates or parties. We consider the case in which there are $n \in \{2, 3\}$ many candidates whose problem is to decide upon which policy to propose (or, equivalently, which position to take in the political spectrum). Each citizen votes for the candidate who has chosen the closest position to her ideal point (because she has single-peaked preferences), and all this is known by the candidates. We assume that the only goal of each candidate is to win the election.⁸ Of course, to complete the specification of the model, we must append to this setting a tie-breaking rule. We postulate in this regard that candidates share equally the votes that they attract together. Each candidate prefers to win the election to a tie for the first place, and the tie for the first place to losing the election.

⁷The political spectrum is of course better modeled as being multidimensional; for instance, in the national elections voters not only care about the position of a candidate on the health care reform but also about his/her position on the tax policy, education reform, social security, and so on. Allowing for multidimensionality in voting models, however, complicates matters to a considerable degree, and hence we choose to confine our attention here to the unidimensional setting.

⁸Once again this is not the most realistic of assumptions. For instance, it would certainly be reasonable to posit that the candidates have policy preferences on their own, and hence, care also about the policy that will be implemented in equilibrium. However, vote maximization is certainly one of the major goals of politicians, and the above model (which is sometimes called the *Downsian* model of political competition) is useful in identifying the implications of such an objective about the pre-election behavior of the candidates. It is by far the most standard in the literature.

Losing the election may or may not be the worst outcome for a candidate, however, depending on whether or not running in the election is costly. We take up each of these possibilities in turn.

Elections when running is not costly

Let us first consider the case where there are two political candidates. In this case the game at hand is a 2-person game with action spaces $A_1 = A_2 = [0, 1]$. The utility function for candidate i is given as

$$u_i(\ell_1, \ell_2) = \begin{cases} 1, & \text{if } i \text{ wins the election alone at the position profile } (\ell_1, \ell_2) \\ \frac{1}{2}, & \text{if there is a tie at the position profile } (\ell_1, \ell_2) \\ 0, & \text{if } i \text{ loses the election at the position profile } (\ell_1, \ell_2), \end{cases}$$

where ℓ_j denotes the policy chosen by candidate $j = 1, 2$.

We wish to find the set of Nash equilibria of this voting game. Rather than computing the best response correspondences of the players, we again launch a direct attack. (You are welcome to verify the validity of the subsequent analysis by using the best response correspondences of the candidates.) Suppose that (ℓ_1, ℓ_2) is a Nash equilibrium. Consider first the possibility that candidate 1 is winning the election at this position profile. Now notice that if $\ell_1 \neq 1/2$, then candidate 2 can force a win by choosing $1/2$. This is because, in this case, she gets all the votes either in $[0, \frac{\ell_1+1/2}{2}]$ (if $\ell_1 > 1/2$) or in $[\frac{\ell_1+1/2}{2}, 1]$ (if $\ell_1 < 1/2$) which add up to more than the half of the total number of votes. But then since ℓ_2 is a best response of candidate 2 against ℓ_1 , it must also guarantee a win for her, contradicting that candidate 1 is the winner in the equilibrium outcome (ℓ_1, ℓ_2) . Therefore, we must have $\ell_1 = 1/2$. But this will not do either, because the best response of candidate 2 against $\ell_1 = 1/2$ is to play $1/2$ which forces a tie, contradicting again the hypothesis that candidate 1 was the winner of the election in equilibrium. Thus, candidate 1 cannot win the election alone in equilibrium, and by symmetry, neither can candidate 2. We thus learn that $\ell_1 = \ell_2$ must be the case, that is, the election is bound to end up in a tie in equilibrium. But it is easily checked that we cannot have $\ell_1 = \ell_2 \neq 1/2$ in equilibrium (either party would then deviate to, say, $1/2$). The only possibility of equilibrium outcome in this game is thus $\ell_1 = \ell_2 = 1/2$, and this is indeed an equilibrium as you can easily verify. The conclusion is that in the unique equilibrium of the game both parties choose the median position.⁹

Life gets more complicated if a third candidate decides to join the race. In the 3-person game that obtains in this case, the action spaces are $A_1 = A_2 = A_3 = [0, 1]$ and the utility

⁹If you are careful, you will notice that the assumption of carefully distributed individuals did not really play a role in arriving at this conclusion. If the distribution is given by an arbitrary continuous density function f on $[0, 1]$ with $f(x) > 0$, the equilibrium would have both parties to locate on the median of this distribution.

function for candidate i is given as

$$u_i(\ell_1, \ell_2, \ell_3) = \begin{cases} 1, & \text{if } i \text{ wins the election alone at the position profile } (\ell_1, \ell_2, \ell_3) \\ \frac{1}{2}, & \text{if } i \text{ there is a tie at the position profile } (\ell_1, \ell_2, \ell_3) \\ 0, & \text{if } i \text{ loses the election at the position profile } (\ell_1, \ell_2, \ell_3), \end{cases}$$

where ℓ_j denotes the policy chosen by candidate $j = 1, 2, 3$. The equilibrium of this game is not a trivial extension of the previous game. Indeed, $(\ell_1, \ell_2, \ell_3) = (1/2, 1/2, 1/2)$ does not correspond to a Nash equilibrium here. For, each candidate is getting approximately the 33% of the total votes at this profile, and by moving slightly to the left (or right) of $1/2$ any of the candidates can increase her share of the votes almost to 50%. None of the candidates is thus playing optimally given the actions of others.

It turns out that there are many Nash equilibria of this 3-person voting game. As an example let us verify that $(\ell_1, \ell_2, \ell_3) = (1/4, 1/4, 3/4)$ is an equilibrium. Begin with observing that candidate 3 wins the election alone at this position profile. Therefore, this candidate obviously does not have any incentive to deviate from $3/4$ given that the other two candidates position themselves at $1/4$. Does candidate 1 (hence candidate 2) has a profitable deviation? No. Given that $(\ell_1, \ell_3) = (1/4, 3/4)$, it is readily observed that if candidate 1 chooses instead of $1/4$ any position in the interval $[0, 1/4]$, then candidate 3 remains as the winner, and if she deviates to any position in the interval $[3/4, 1]$, the candidate 2 becomes the winner alone. Less clear is the implication of choosing a policy in the interval $(1/4, 3/4)$. The key observation here is that by doing so candidate 1 would attract the votes that belong to $\left[\frac{\ell_1+1/4}{2}, \frac{\ell_1+3/4}{2}\right]$. Thus in this case candidate 1 would get exactly the 25% of the total vote (see Figure 5.) But either candidate 2 ($3/4 > \ell_1 \geq 1/2$) or candidate 3 (if $1/2 \geq \ell_1 > 1/4$) is bound to collect 37.5% of the votes in this case. Therefore, choosing $1/4$ is as good as choosing any other position in $[0, 1]$ for candidate 1 given the actions of others, she maintains a payoff level of 0 with any such choice. So, at the profile $(\ell_1, \ell_2, \ell_3) = (1/4, 1/4, 3/4)$, neither candidate 1 nor candidate 2 can force a win by means of a unilateral deviation, and we conclude that this outcome is a Nash equilibrium. (*Challenge*: Compute all the Nash equilibria of this game.)

But in voting problems the issue of coalitions arise very naturally. So we better ask if the equilibrium $(1/4, 1/4, 3/4)$ is actually strong or not. Indeed, it is not. For, candidates 1 and 2 can jointly deviate at this profile to, say, $3/4 - \varepsilon$ for small $\varepsilon > 0$, and thus force a win (which yields a payoff of $1/2$ to each). What is more, there is no strong Nash equilibrium of this game. (*Challenge*: Prove this.)

Elections when running is costly

In this case, staying out is a meaningful alternative for each political candidate. Consequently, we model the situation as a game in strategic form by setting, for each candidate i ,

$A_i = [0, 1] \cup \{\text{stay out}\}$ and

$$u_i(\ell) = \begin{cases} 1, & \text{if } i \text{ wins the election alone at the position profile } \ell \\ \frac{1}{2}, & \text{if } i \text{ ties for the first place at the position profile } \ell \\ 0, & \text{if } i \text{ stays out at the position profile } \ell \\ -1, & \text{if } i \text{ runs but loses the election at the position profile } \ell \end{cases}$$

where $\ell \in \prod_{i=1}^n A_i$ with $n \in \{2, 3\}$.

The equilibrium analysis of this game is essentially identical to the first game considered above when $n = 2$. Consequently, we leave the related analysis as an

Exercise. Prove: If $n = 2$, the unique Nash equilibrium of the game defined above is $(1/2, 1/2)$.

Once again, life is more complicated in the 3-person scenario, but now this is not because of the multiplicity of equilibria. On the contrary, this game has no Nash equilibrium when $n = 3$. A sketch of proof can be given as follows. First observe that, since each candidate can avoid losing by staying out of the election, all running candidates must tie for the first place in any equilibrium. Moreover, there cannot be only one running candidate in equilibrium, for otherwise, any other player may choose the same location with the running candidate and forces a tie for the first place (which is better than staying out). Similarly, it cannot be that everyone stays out in equilibrium. Therefore, in any given equilibrium, there must exist two or more running candidates who tie for the first place. Consider first the possibility that there are exactly two such candidates. Then, by the exercise above, both candidates must be choosing $1/2$. But since the running candidates share the total votes, the remaining candidate can force a win by choosing slightly to the left (or right) of $1/2$. Thus staying out cannot be a best response for this candidate, contradicting that we are at an equilibrium. The final possibility is the case in which all three candidates choose not to stay out and tie for the first place. Suppose that (ℓ_1, ℓ_2, ℓ_3) is such an equilibrium. If $\ell_1 = \ell_2 = \ell_3$, any one of the candidates can profitably deviate and force a win (why?), so at least two components of (ℓ_1, ℓ_2, ℓ_3) must be distinct. Suppose that $\ell_1 \neq \ell_2 = \ell_3$. In this case, candidate 1 can force a win by getting very close to ℓ_2 (see this?), and hence she cannot be best responding in the profile (ℓ_1, ℓ_2, ℓ_3) . The two other possibilities in which exactly two components of (ℓ_1, ℓ_2, ℓ_3) are distinct are ruled out similarly. We are then left with the following final possibility: $\ell_1 \neq \ell_2 \neq \ell_3 \neq \ell_1$. To rule out this case as well, we pick the leftist candidate, call her i (so we have $\ell_i = \min\{\ell_1, \ell_2, \ell_3\}$), and observe that this candidate can force a win by choosing a position very close to the median of $\{\ell_1, \ell_2, \ell_3\}$. So, finally, we can conclude that there does not exist a Nash equilibrium for the 3-person voting game at hand.

The following table summarizes our findings in the four spatial voting games we have

examined above.

		The number of candidates	
		2	3
Running is	costly	the only equilibrium is the median position	a Nash equilibrium does not exist
	not costly	the only equilibrium is the median position	there are many Nash equilibria but a strong Nash equilibrium does not exist

It is illuminating to observe how seemingly minor alterations in these voting models result in such vast changes in the set of equilibria.

Mixed Strategy Equilibrium

Levent Koçkesen

1 Introduction

Up to now we have assumed that the only choice available to players was to pick an action from the set of available actions. In some situations a player may want to randomize between several actions. If a player chooses which action to play randomly, we say that the player is using a *mixed strategy*, as opposed to a *pure strategy*. In a pure strategy the player chooses an action for sure, whereas in a mixed strategy, she chooses a probability distribution over the set of actions available to her. In this section we will analyze the implications of allowing players to use mixed strategies.

As a simple illustration, consider the following matching-pennies game.

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

If we restrict players' strategies only to actions, as we have done so far, this game has no Nash equilibrium (check), i.e., it has no Nash equilibrium in pure strategies. Since we have argued that Nash equilibrium is a necessary condition for a steady state, does that mean that the matching-pennies game has no steady state? To answer this question let us allow players to use mixed strategies. In particular, let each player play H and T with half probability each. We claim that this choice of strategies constitute a steady-state, in the sense that if each player predicts that the other player will play in this manner, then she has no reason not to play in the specified manner. Since player 2 plays H with probability $1/2$, the expected payoff of player 1 if she plays H is $(1/2)(1) + (1/2)(-1) = 0$. Similarly, the expected payoff to action T is 0. Therefore, player 1 has no reason to deviate from playing H and T with probability $1/2$ each. Similarly, if player 2 predicts that player 1 will play H and T with half probability each, she has no reason to deviate from doing the same. This shows that the strategy profile where player 1 and 2 play H and T with half probability each is a steady-state of this situation. We say that playing H and T with probabilities $1/2$ and $1/2$ respectively constitutes a mixed strategy equilibrium of this game.

If we assume that players repeatedly play this game and forecast each other's action on the basis of past play, then each player actually has an incentive to adopt a mixed strategy with these probabilities. If, for example, player 1 plays H constantly, rather than the above mixed

strategy, then it is reasonable that player 2 will come to expect him to play H again and play her best response, which is T . This will result in player 1 getting -1 as long as he continues playing H . Therefore, he should try to be unpredictable, for as soon as his opponent becomes able to predict his action, she will be able to take advantage of the situation. Therefore, player 1 should try to mimic playing a mixed strategy by playing H and T with frequencies $1/2$ and $1/2$.

Consider the Hawk-Dove game for a another motivation.

	H	D
H	0, 0	6, 1
D	1, 6	3, 3

Suppose each period two randomly selected individuals, who both belong to a large population, play this game. Also suppose that $3/4$ of the population plays H (is hawkish) and $1/4$ plays D (is dovish), but no player can identify the opponent's type before the game is played. We claim that this is a stable population composition. Since the opponent is chosen randomly from a large population, each player expects the opponent to play H with probability $3/4$ and D with probability $1/4$. Would a dovish player do better if she were a hawkish player? Well, on average a dovish player gets a payoff of $(3/4)(1) + (1/4)(3) = 3/2$. A hawkish player gets $(3/4)(0) + (1/4)(6) = 3/2$ as well. Therefore, neither type of player has a reason to change his behavior.

2 Mixed Strategies and Expected Payoffs

Definition 1 A *mixed strategy* α_i for player i , is a probability distribution over his set of available actions, A_i . In other words, if player i has m actions available, a mixed strategy is an m dimensional vector $(\alpha_i^1, \alpha_i^2, \dots, \alpha_i^m)$ such that $\alpha_i^k \geq 0$, for all $k = 1, 2, \dots, m$, and $\sum_{k=1}^m \alpha_i^k = 1$.

We will denote by $\alpha_i(a_i)$ the probability assigned to action a_i by the mixed strategy α_i . Let $\Delta(X)$ denote the set of all probability distributions on a set X . Then, any mixed strategy α_i for player i is an element of $\Delta(A_i)$, i.e., $\alpha_i \in \Delta(A_i)$. Following the convention we developed for action profiles, we will denote by $\alpha = (\alpha_i)_{i \in N}$ a mixed strategy profile, i.e., a mixed strategy for each player in the game. To denote the strategy profile in which player i plays α'_i and the rest of the players play α_j^* , $j \neq i$, we will use $(\alpha'_i, \alpha_{-i}^*)$. Unless otherwise stated, we will assume that players choose their mixed strategies independently.

Notice that not all actions have to receive a positive probability in a mixed strategy. Therefore, it is also possible to see pure strategies as degenerate mixed strategies, in which all but one action is played with zero probability.

Let us illustrate these concepts by using the Battle of the Sexes game that we introduced before:

	m	o
m	2,1	0,0
o	0,0	1,2

A possible mixed strategy for player 1 is $(1/2, 1/2)$, or $\alpha_1(m) = \alpha_1(o) = 1/2$. Another is $(1/3, 2/3)$, or $\alpha_1(m) = 1/3, \alpha_1(o) = 2/3$. For player 2, we may have $(2/3, 1/3)$, i.e., $\alpha_2(m) = 2/3, \alpha_2(o) = 1/3$, as a possible mixed strategy. A mixed strategy profile could be $((1/2, 1/2), (2/3, 1/3))$ another could be $((1/3, 2/3), (2/3, 1/3))$. Notice that we always have $\alpha_1(o) = 1 - \alpha_1(m)$ and $\alpha_2(o) = 1 - \alpha_2(m)$ simply because probabilities have to add up to one. Therefore, sometimes we may want to simplify the notation by defining, say $p \equiv \alpha_1(m)$, $q \equiv \alpha_2(m)$, and using (p, q) to denote a strategy profile, where player 1 chooses m with probability p and action o with probability $1 - p$, and player 2 chooses m with probability q and action o with probability $1 - q$. Notice, if there were 3 actions for a player, then we would need at least two numbers to specify a mixed strategy for that player.

Once we allow players to use mixed strategies, the outcomes are not deterministic anymore. For example if both players play m with probability $1/2$ in the BoS game, then each action profile is obtained with probability $1/4$. Therefore, we have to specify players' preferences over lotteries, i.e., over probability distributions over outcomes, rather than preferences over certain outcomes. We will assume that players' preferences satisfy the assumptions of **Von Neumann and Morgenstern** so that the payoff to an uncertain outcome is the weighted average of the payoffs to underlying certain outcomes, weight attached to each outcome being the probability with which that outcome occurs. (See Dutta, P., ch. 27 for more on this). In other words, we assume that for each player i , there is a payoff function u_i defined over the certain outcomes $a \in A$, such that the player's preferences over lotteries on A can be represented by the expected value of u_i . If each outcome $a \in A$ occurs with probability $p(a)$, then the expected payoff of player i is

$$u_i(p) \equiv \sum_{a \in A} p(a) u_i(a).$$

Example 2 For example, in the BoS game if each player i plays the mixed strategy α_i , then the expected payoff of player i is given by

$$\begin{aligned} u_1(\alpha_1, \alpha_2) &= \alpha_1(m) \alpha_2(m) u_i(m, m) + \alpha_1(m) \alpha_2(o) u_i(m, o) \\ &\quad + \alpha_1(o) \alpha_2(m) u_i(o, m) + \alpha_1(o) \alpha_2(o) u_i(o, o) \\ &= \alpha_1(m) [\alpha_2(m) u_i(m, m) + \alpha_2(o) u_i(m, o)] \\ &\quad + \alpha_1(o) [\alpha_2(m) u_i(o, m) + \alpha_2(o) u_i(o, o)] \\ &= \alpha_1(m) u_1(m, \alpha_2) + \alpha_1(o) u_1(o, \alpha_2), \end{aligned}$$

or,

$$\begin{aligned} u_1(\alpha_1, \alpha_2) &= \alpha_1(m) \alpha_2(m) (2) + \alpha_1(m) \alpha_2(o) u_i(m, o) (0) \\ &\quad + \alpha_1(o) \alpha_2(m) (0) + \alpha_1(o) \alpha_2(o) (1) \\ &= 2\alpha_1(m) \alpha_2(m) + \alpha_1(o) \alpha_2(o), \end{aligned}$$

and that of player 2 is

$$u_2(\alpha_1, \alpha_2) = \alpha_1(m) \alpha_2(m) + 2\alpha_1(o) \alpha_2(o).$$

Notice that, since $\alpha_i(o) = 1 - \alpha_i(m)$, we can write these expected payoffs as

$$\begin{aligned} u_1(\alpha_1, \alpha_2) &= 2\alpha_1(m) \alpha_2(m) + (1 - \alpha_1(m))(1 - \alpha_2(m)) \\ &= 1 - \alpha_2(m) + \alpha_1(m) [3\alpha_2(m) - 1] \end{aligned}$$

and

$$u_2(\alpha_1, \alpha_2) = 2 - 2\alpha_1(m) + \alpha_2(m) [3\alpha_1(m) - 2].$$

For example, if player 1 plays m for sure, i.e., $\alpha_1(m) = 1$, and player 2 plays m with probability $1/3$, then

$$\begin{aligned} u_1(\alpha_1, \alpha_2) &= 1 - 1/3 + 1 [3 \times (1/3) - 1] \\ &= 2/3 \end{aligned}$$

and

$$\begin{aligned} u_2(\alpha_1, \alpha_2) &= 2 - 2(1) + (1/3) [3(1) - 2] \\ &= 1/3. \end{aligned}$$

Definition 3 The **support** of a mixed strategy α_i is the set of actions to which α_i assigns a positive probability, i.e.,

$$\text{supp}(\alpha_i) = \{a_i \in A_i : \alpha_i(a_i) > 0\}.$$

In the above example we have, $\text{supp}(\alpha_1) = \{m\}$, and $\text{supp}(\alpha_2) = \{m, o\}$.

3 Mixed Strategy Equilibrium

Definition 4 **Best response correspondence** of player i is the set of mixed strategies which are optimal given the other players' mixed strategies. In other words:

$$B_i(\alpha_{-i}) = \arg \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \alpha_{-i}).$$

Example 5 Suppose $\alpha_2(m) = 1/2$. Then, we have

$$\begin{aligned} u_1(\alpha_1, (1/2, 1/2)) &= 1 - 1/2 + \alpha_1(m) [3(1/2) - 1] \\ &= \frac{1}{2} + \frac{1}{2}\alpha_1(m) \end{aligned}$$

therefore,

$$B_1((1/2, 1/2)) = \{(1, 0)\}.$$

In general, letting $p \equiv \alpha_1(m)$, and $q \equiv \alpha_2(m)$, we can express the best response of player 1 in terms of optimal choice of p in response to q

$$B_1(q) = \begin{cases} \{1\}, & \text{if } q > 1/3 \\ [0, 1], & \text{if } q = 1/3 \\ \{0\}, & \text{if } q < 1/3 \end{cases}.$$

The best response correspondence of player 2, i.e., optimal choices of q in response to p , is

$$B_2(p) = \begin{cases} \{1\}, & \text{if } p > 2/3 \\ [0, 1], & \text{if } p = 2/3 \\ \{0\}, & \text{if } p < 2/3 \end{cases}.$$

[See Figure 1.]

Definition 6 A mixed strategy equilibrium is a mixed strategy profile $(\alpha_1^*, \dots, \alpha_n^*)$ such that, for all $i = 1, \dots, n$

$$\alpha_i^* \in \arg \max_{\alpha_i \in \Delta(A_i)} u_i(\alpha_i, \alpha_{-i}^*)$$

or

$$\alpha_i^* \in B_i(\alpha_{-i}^*).$$

In the Battle of the Sexes game, then, the set of mixed strategy Nash equilibria is

$$\{((1, 0), (0, 1)), ((0, 1), (1, 0)), ((2/3, 1/3), (1/3, 2/3))\}.$$

Alternatively, we may say that the set of mixed strategy equilibria is

$$\{(\alpha_1(m), \alpha_2(m)) : (1, 0), (0, 1), (2/3, 1/3)\}.$$

Remark 1 A mixed strategy α_i is a best response to α_{-i} if and only if every action in the support of α_i is itself a best response to α_{-i} . Otherwise, player i could transfer probability from the action which is not a best response to an action which is a best response and strictly increase his payoff.




Figure 1: Best Response Correspondences in BoS game

Remark 2 *This suggests an easy way to find mixed strategy Nash equilibrium. A mixed strategy profile α^* is a mixed strategy Nash equilibrium if and only if for each player i , each action in the support of α_i^* is a best response to α_{-i}^* . In other words, each action in the support of α_i^* yields the same expected payoff when played against α_{-i}^* , and no other action yields a strictly higher payoff.*

Remark 3 *One implication of the above remark is that a nondegenerate mixed strategy equilibrium is not strict.*

Example 7 *In the BoS game, if (α_1^*, α_2^*) is a mixed strategy equilibrium with $\text{supp}(\alpha_1^*) = \text{supp}(\alpha_2^*) = \{m, o\}$, then it must be that the expected payoffs to m and o are the same for both player against α_{-i}^* . In other words, for player 1*

$$2\alpha_2^*(m) = 1 - \alpha_2^*(m)$$

and for player 2

$$\alpha_1^*(m) = 2 - 2\alpha_1^*(m)$$

which imply that

$$\begin{aligned}\alpha_2^*(m) &= 1/3 \\ \alpha_1^*(m) &= 2/3\end{aligned}$$

Proposition 8 *Every finite strategic form game has a mixed strategy equilibrium.*

4 Dominated Actions and Mixed Strategies

In earlier lectures we defined an action to be weakly or strictly dominated, only if there existed another action which weakly or strictly dominated that action. However, it is possible that an action is not dominated by any other action, yet it is dominated by a mixed strategy.

Definition 9 *In a strategic form game, player i 's mixed strategy α_i^* strictly dominates her action a'_i if*

$$u_i(\alpha_i, a_{-i}) > u_i(a'_i, a_{-i}) \text{ for all } a_{-i} \in A_{-i}.$$

Example 10 *Consider the following game,*

	<i>L</i>	<i>R</i>
<i>T</i>	1, 1	1, 0
<i>M</i>	3, 0	0, 3
<i>B</i>	0, 1	4, 1

Clearly, no action dominates T , but the mixed strategy $\alpha_1(M) = 1/2$, $\alpha_1(B) = 1/2$ strictly dominates T .

Remark 4 *A strictly dominated action is never used with positive probability in a mixed strategy equilibrium*

To find the mixed strategy equilibria in games where one of the players have more than two actions one should first look for strictly dominated actions and eliminate them. (see example 118.2 in Osborne chapter 4 on my web page).

Bayesian Games

Levent Koçkesen

So far we have assumed that all players had perfect information regarding the elements of a game. These are called games with complete information. A game with **incomplete information**, on the other hand, tries to model situations in which some players have private information before the game begins. The initial private information is called the **type** of the player. For example, types could be the privately observed costs in an oligopoly game, or privately known valuations of an object in an auction, etc.

1 Preliminaries

A **Bayesian game** is a strategic form game with incomplete information. It consists of:

- a set of players, $N = \{1, \dots, n\}$,
and for each $i \in N$
- an action set, A_i , ($A = \times_{i \in N} A_i$)
- a type set, Θ_i , ($\Theta = \times_{i \in N} \Theta_i$)
- a probability function,

$$p_i : \Theta_i \rightarrow \Delta(\Theta_{-i})$$

- a payoff function,

$$u_i : A \times \Theta \rightarrow \mathbf{R}.$$

The function p_i summarizes what player i believes about the types of the other players given her type. So, $p_i(\theta_{-i}|\theta_i)$ is the conditional probability assigned to the type profile $\theta_{-i} \in \Theta_{-i}$. Similarly, $u_i(a|\theta)$ is the payoff of player i when the action profile is a and the type profile is θ .

We call a Bayesian game **finite** if N , A_i and Θ_i are all finite, for all $i \in N$. A **pure strategy** for player i in a Bayesian game is a function which maps player i 's type into her action set

$$a_i : \Theta_i \rightarrow A_i,$$

so that $a_i(\theta_i)$ is the action choice of type θ_i of player i .

A mixed strategy for player i is

$$\alpha_i : \Theta_i \rightarrow \Delta(A_i)$$

so that $\alpha_i(a_i|\theta_i)$ is the probability assigned by α_i to action a_i by type θ_i of player i .

Suppose there are two players, player 1 and 2 and for each player there are two possible types. Player i 's possible types are θ_i and θ'_i . Suppose that the types are independently distributed and the probability of θ_1 is p and the probability of θ_2 is q . For a given pure strategy profile a^* the expected payoff of player 1 of type θ_1 is

$$qu_1(a_1^*(\theta_1), a_2^*(\theta_2) | \theta_1, \theta_2) + (1 - q) u_1(a_1^*(\theta_1), a_2^*(\theta'_2) | \theta_1, \theta'_2).$$

Similarly, for a given mixed strategy profile α^* the expected payoff of player 1 of type θ_1 is

$$q \sum_{a \in A} \alpha_1^*(a_1 | \theta_1) \alpha_2^*(a_2 | \theta_2) u_1(a_1, a_2 | \theta_1, \theta_2) + (1 - q) \sum_{a \in A} \alpha_1^*(a_1 | \theta_1) \alpha_2^*(a_2 | \theta'_2) u_1(a_1, a_2 | \theta_1, \theta'_2)$$

2 Bayesian Equilibrium

Definition 1 A *Bayesian equilibrium* of a Bayesian game is a mixed strategy profile $\alpha = (\alpha_i)_{i \in N}$, such that for every player $i \in N$ and every type $\theta_i \in \Theta_i$, we have

$$\alpha_i(\cdot | \theta_i) \in \arg \max_{\gamma \in \Delta(A_i)} \sum_{\theta_{-i} \in \Theta_{-i}} p_i(\theta_{-i} | \theta_i) \sum_{a \in A} \left(\prod_{j \in N \setminus \{i\}} \alpha_j(a_j | \theta_j) \right) \gamma(a_i) u_i(a | \theta).$$

Remark 1 Type, in general, can be any private information that is relevant to the player's decision making, such as the payoff function, player's beliefs about other players' payoff functions, her beliefs about what other players believe her beliefs are, and so on.

Remark 2 Notice that, in the definition of a Bayesian equilibrium we need to specify strategies for each type of a player, even if in the actual game that is played all but one of these types are non-existent. This is because, given a player's incomplete information, analysis of that player's decision problem requires us to consider what each type of the other players would do, if they were to play the game.

3 Some Examples

3.1 Battle of the Sexes with incomplete information.

Suppose player 2 has perfect information and two types l and h . Type l loves going out with player 1 whereas type h hates it. Player 1 has only one type and does not know which type is player 2. Her beliefs place probability 1/2 on each type. The following tables give the payoffs to each action and type profile:

	B	S		B	S
B	2,1	0,0	B	2,0	0,2
S	0,0	1,2	S	0,1	1,0
	<i>type l</i>			<i>type h</i>	

We can represent this situation as a Bayesian game:

- $N = \{1, 2\}$
- $A_1 = A_2 = \{B, S\}$
- $\Theta_1 = \{x\}, \Theta_2 = \{l, h\}$
- $p_1(l|x) = p_1(h|x) = 1/2, p_2(x|l) = p_2(x|h) = 1.$
- u_1, u_2 are given in the tables above.

Since player 1 has only one type (i.e., his type is common knowledge) we will omit references to his type from now on.

Let us find the Bayesian equilibria of this game by analyzing the decision problem of each player of each type:

Player 2 of type l : Given player 1's strategy α_1 , his expected payoff to

- action B is $\alpha_1(B)$,
- action S is $2(1 - \alpha_1(B))$

so that his best response is to play B if $\alpha_1(B) > 2/3$ and to play S if $\alpha_1(B) < 2/3$.

Player 2 of type h : Given player 1's strategy α_1 , his expected payoff to

- action B is $(1 - \alpha_1(B))$,
- action S is $2\alpha_1(B)$

so that his best response is to play B if $\alpha_1(B) < 1/3$ and to play S if $\alpha_1(B) > 1/3$.

Player 1: Given player 2's strategy $\alpha_2(.|l)$ and $\alpha_2(.|h)$, her expected payoff to

- action B is

$$\frac{1}{2}\alpha_2(B|l)(2) + \frac{1}{2}\alpha_2(B|h)(2) = \alpha_2(B|l) + \alpha_2(B|h),$$

- action S is

$$\frac{1}{2}(1 - \alpha_2(B|l))(1) + \frac{1}{2}(1 - \alpha_2(B|h))(1) = 1 - \frac{\alpha_2(B|l) + \alpha_2(B|h)}{2}.$$

Therefore, her best response is to play B if $\alpha_2(B|l) + \alpha_2(B|h) > 2/3$ and to play S if $\alpha_2(B|l) + \alpha_2(B|h) < 2/3$.

Let us first check if there is a pure strategy equilibrium in which both types of player 2 play B , i.e. $\alpha_2(B|l) = \alpha_2(B|h) = 1$. In this case player 1's best response is to play B as well to which playing B is not a best response for player 2 type h . Similarly check that $\alpha_2(B|l) = \alpha_2(B|h) = 0$ and $\alpha_2(B|l) = 0$ and $\alpha_2(B|h) = 1$ cannot be part of a Bayesian equilibrium. Let's check if $\alpha_2(B|l) = 1$ and $\alpha_2(B|h) = 0$ could be part of an equilibrium. In this case player 1's best response is to play B . Player 2 type l 's best response is to play B and that of type h is S . Therefore,

$$(\alpha_1(B|x), \alpha_2(B|l), \alpha_2(B|h)) = (1, 1, 0)$$

is a Bayesian equilibrium.

Clearly, there is no equilibrium in which both types of player 2 mixes. Suppose only type l mixes. Then, $\alpha_1(B) = 2/3$, which implies that $\alpha_2(B|l) + \alpha_2(B|h) = 2/3$. This, in turn, implies that $\alpha_2(B|h) = 0$. Since $\alpha_2(B|h) = 0$ is a best response to $\alpha_1(B|x) = 2/3$, the following is another Bayesian equilibrium of this game

$$(\alpha_1(B), \alpha_2(B|l), \alpha_2(B|h)) = (2/3, 2/3, 0).$$

As an exercise show there is one more equilibrium given by

$$(\alpha_1(B), \alpha_2(B|l), \alpha_2(B|h)) = (1/3, 0, 2/3).$$

3.2 Cournot Duopoly with incomplete information.

The profit functions are given by

$$u_i = q_i(\theta_i - q_i - q_j).$$

Firm 1 has one type $\theta_1 = 1$, but firm 2 has private information about its type θ_2 . Firm 1 believes that $\theta_2 = 3/4$ with probability $1/2$ and $\theta_2 = 5/4$ with probability $1/2$, and this belief is common knowledge.

We will look for a pure strategy equilibrium of this game. Firm 2 of type θ_2 's decision problem is to

$$\max_{q_2} q_2(\theta_2 - q_1 - q_2)$$

which is solved at

$$q_2^*(\theta_2) = \frac{\theta_2 - q_1}{2}.$$

Firm 1's decision problem, on the other hand, is

$$\max_{q_1} \left\{ \frac{1}{2} q_1 (1 - q_1 - q_2^*(3/4)) + \frac{1}{2} q_1 (1 - q_1 - q_2^*(5/4)) \right\}$$

which is solved at

$$q_1^* = \frac{2 - q_2^*(3/4) - q_2^*(5/4)}{4}.$$

Solving yields,

$$q_1^* = \frac{1}{3}, q_2^*(3/4) = \frac{11}{24}, q_2^*(5/4) = \frac{5}{24}.$$

Auctions

Many economic transactions are conducted through auctions. Governments sell treasury bills, foreign exchange, publicly owned companies, mineral rights, and more recently airwave spectrum rights via auctions. Art work, antiques, cars, and houses are also sold by auctions. Government contracts are awarded by procurement auctions, which are also used by firms to buy inputs or to subcontract work. Takeover battles are effectively auctions as well and auction theory has been applied to areas as diverse as queues, wars of attrition, and lobbying contests.¹

There are four commonly used and studied forms of auctions:

1. *ascending-bid auction* (also called the open, oral, or, English auction): the price is raised until only one bidder remains, and that bidder wins the object at the final price.
2. *descending-bid auction* (also called Dutch auction): the auctioneer starts at a very high price and lowers it continuously until someone accepts the currently announced price. That bidder wins the object at that price.
3. *first-price sealed bid auction*: each bidder submits her bid in a sealed envelope without seeing others' bids, and the object is sold to the highest bidder at her bid.
4. *second-price sealed bid auction* (also known as Vickrey auction²). Bidders submit their bids in a sealed envelope, the highest bidder wins but pays the second highest bid.

Auctions also differ with respect to the valuation of the bidders. In a *private value auction* each bidder's valuation is known only by the bidder, as it could be the case, for example, in an artwork or antique auction. In a *common value auction*, the actual value of the object is the same for everyone, but bidders have different private information regarding that value. For example, the value of an oil tract or a company maybe the same for everybody but different bidders may have different estimates of that value.

We will analyze sealed bid auctions, not only because they are simpler to analyze but also because in the private values case, the first-price sealed bid auction is strategically equivalent to descending bid auction and the second-price sealed bid auction is strategically equivalent to ascending bid auction.

¹For a good introductory survey to the auction theory see Paul Klemperer (1999), "Auction Theory: A Guide to the Literature," *Journal of Economic Surveys*, 13(3), July 1999, pp. 227-286.

²Named after William Vickrey of Columbia University who was awarded the Nobel Prize in economics in 1996.

1 Independent Private Values

Previously, we have looked at two forms of auctions, namely First Price and Second Price Auctions, in a complete information framework in which each bidder knew the valuations of every other bidder. In this section we relax the complete information assumption and revisit these two form of auctions. In particular, we will assume that each bidder knows only her own valuation, and the valuations are independently distributed random variables whose distributions are common knowledge.

The following elements define the general form of an auction that we will analyze:

- Set of bidders, $N = \{1, 2, \dots, n\}$,
and for each $i \in N$
- a type set (set of possible valuations), $\Theta_i = [\underline{v}, \bar{v}]$, $\underline{v} \geq 0$.
- an action set, $A_i = \mathbf{R}_+$ (actions are bids)
- a belief function: player i believes that her opponents' valuations are independent draws from a distribution function F that is strictly increasing and continuous on $[\underline{v}, \bar{v}]$.
- a payoff function, which is defined for any $a \in A$, $v \in \Theta$ as follows

$$u_i(a, v) = \begin{cases} \frac{v_i - P(a)}{m}, & \text{if } a_j \leq a_i \text{ for all } j \neq i, \text{ and } |\{j : a_j = a_i\}| = m \\ 0, & \text{if } a_j > a_i \text{ for some } j \neq i \end{cases}$$

where $P(a)$ is the price paid by the winner if the bid profile is a . Notice that in the case of a tie the object is divided equally among all winners.

1.1 Second Price Auctions

In this design, highest bidder wins and pays a price equal to the second highest bid. Although there are many Bayesian equilibria of second price auctions, bidding own valuation v_i is weakly dominant for each player i . To see this let x be the highest of the other bids and consider bidding $a'_i < v_i$, v_i , and $a''_i > v_i$. Depending upon the value of x , the following table gives the payoffs to each of these actions by i

	$x \leq a'_i$	$a'_i < x < v_i$	$x = v_i$	$v_i < x \leq a''_i$	$a''_i < x$
a'_i	win/tie; pay x	lose	lose	lose	lose
v_i	win; pay x	win; pay x	tie; pay v_i	lose	lose
a''_i	win; pay x	win; pay x	win; pay v_i	win/tie; pay x	lose

By bidding smaller than v_i , you sometimes lose when you should win ($a_i < x < v_i$) and by bidding more than v_i , you sometimes win when you should lose ($a_i > x > v_i$).

1.2 First Price Auctions

In first price auctions, the highest bidder wins and pays her bid. Let us denote the bid of player with type v_i by $\beta_i(v_i)$ and look for symmetric equilibria, i.e. $\beta_i(v) = \beta(v)$ for all $i \in N$. First, although we will not attempt to do so here, it is possible to show that strategies $\beta_i(v_i)$, and hence $\beta(v)$, are strictly increasing and continuous on $[\underline{v}, \bar{v}]$. (see Fudenberg and Tirole, 1991). So, let's assume that they are, and check if they are once we locate a possible equilibrium.

The expected payoff of player with type v who bids b when all the others are bidding according to β is given by

$$\begin{aligned}(v - b) \text{prob}(\text{highest bid is } b) &= (v - b) (\text{prob}(\beta(v) \leq b))^{n-1} \\ &= (v - b) F(\beta^{-1}(b))^{n-1}.\end{aligned}$$

(Because v_i are independently distributed). The first order condition for maximizing the expected payoff is

$$-F(\beta^{-1}(b))^{n-1} + (v - b)(n - 1)F(\beta^{-1}(b))^{n-2}F'(\beta^{-1}(b))\frac{1}{\beta'(\beta^{-1}(b))} = 0.$$

by the fact that β is almost everywhere differentiable (since it is strictly increasing), and by the inverse function theorem. For $\beta(v)$ to be an equilibrium first order condition must hold when we substitute $\beta(v)$ for b ,

$$-F(v)^{n-1} + (v - \beta(v))(n - 1)F(v)^{n-2}F'(v)\frac{1}{\beta'(v)} = 0,$$

or

$$\beta'(v)F(v)^{n-1} + (n - 1)\beta(v)F'(v)F(v)^{n-2} = (n - 1)vF'(v)F(v)^{n-2}$$

which is a differential equation in β . Integrating both sides, we get

$$\begin{aligned}\beta(v)F(v)^{n-1} &= \int_{\underline{v}}^v (n - 1)xF(x)^{n-2}F'(x)dx \\ &= vF(v)^{n-1} - \int_{\underline{v}}^v F(x)^{n-1}dx.\end{aligned}$$

Solving for $\beta(v)$,

$$\beta(v) = v - \frac{\int_{\underline{v}}^v F(x)^{n-1}dx}{F(v)^{n-1}}.$$

One can easily show that β is continuous and strictly increasing in v as we hypothesized. Furthermore, notice that $\beta(\underline{v}) = \underline{v}$, but $\beta(v) < v$ for $v > \underline{v}$. That is, except the player with the lowest valuation, everybody bids less than her valuation. As an exercise, let's calculate β assuming F is uniform on $[0, 1]$, i.e., $F(x) = x$.

$$\beta(v) = v - \frac{\int_0^v x^{n-1} dx}{v^{n-1}} = v - \frac{1}{v^{n-1}} \frac{v^n}{n} = \frac{n-1}{n} v.$$

Uniform example solved explicitly: Let's look for a symmetric equilibrium of the form $\beta(v) = av$. The expected payoff of player with type v who bids b when all the others are bidding according to β is given by

$$\begin{aligned} (v-b) \text{prob}(\text{highest bid is } b) &= (v-b) (\text{prob}(av \leq b))^{n-1} \\ &= (v-b) (b/a)^{n-1}. \end{aligned}$$

The first order condition for maximizing the expected payoff is

$$(v-b)(n-1) = b,$$

which is solved at

$$b = \frac{n-1}{n} v.$$

1.3 All-Pay Auctions

Consider an auction in which the highest bidder wins the auction but every bidder pays his/her bid. This model could model bribes, political contests, Olympic competition, war-of-attritions, etc. Again, let's look for a symmetric equilibrium, $\beta_i(v) = \beta(v)$ for all $i \in N$. The expected payoff of player with type v who bids b when all the others are bidding according to β is given by

$$\begin{aligned} v \times \text{prob}(\text{highest bid is } b) - b &= v \times \text{prob}(\beta(v) \leq b)^{n-1} - b \\ &= v F(\beta^{-1}(b))^{n-1} - b. \end{aligned}$$

Let F be uniform over $[0, 1]$. Then, this becomes

$$v (\beta^{-1}(b))^{n-1} - b.$$

The first order condition for maximizing the expected payoff is

$$-1 + v(n-1) (\beta^{-1}(b))^{n-2} \frac{1}{\beta'(\beta^{-1}(b))} = 0.$$

For $\beta(v)$ to be an equilibrium first order condition must hold when we substitute $\beta(v)$ for b ,

$$-1 + v(n-1) v^{n-2} \frac{1}{\beta'(v)} = 0,$$

or

$$\beta'(v) = (n-1) v^{n-1}$$

which is a differential equation in β . Integrating both sides, we get

$$\begin{aligned}\beta(v) &= \int_0^v (n-1) x^{n-1} dx \\ &= \frac{(n-1)}{n} v^n.\end{aligned}$$

Notice that as n increases the equilibrium bid decreases.

2 Revenue Equivalence

In second price auctions each bidder bids her value and pays the second highest. Therefore, the expected revenue of the seller is the expected second highest value. In a first price auction, the highest bidder is the one with the highest value and bids a function of her value, which is $\frac{n-1}{n} v_{\max}$ in our example above. Therefore, the seller's expected revenue in a first price and a second price auction depends on the expectation of the highest and the second highest value, respectively. Given that there are n bidders who each has a value (drawn independently from a common distribution), what are the expected values of the highest and second highest values? Order statistics provide the answer.

Order Statistics

Suppose that v is a real-valued random variable with distribution function F and density function f . Also suppose that n independent values are drawn from the same distribution to form a random sample $(v_1, v_2, \dots, v_n)$. Let $v_{(k)}$ denote the k th smallest of $(v_1, v_2, \dots, v_n)$ and call it k th *order statistic*. In particular $v_{(n)}$ is the highest and $v_{(n-1)}$ is the second highest order statistics. Let F_k denote the distribution function of $v_{(k)}$. Let's start with the distribution function of $v_{(n)}$.

$$\begin{aligned}F_n(x) &= \text{prob}(v_{(n)} \leq x) = \text{prob}(\text{all } v_i \leq x) \\ &= [F(x)]^n.\end{aligned}$$

Similarly,

$$\begin{aligned}F_{n-1}(x) &= \text{prob}(v_{(n-1)} \leq x) \\ &= \text{prob}(\text{either } n \text{ or } (n-1) \text{ of } v\text{'s are } \leq x) \\ &= [F(x)]^n + n(1 - F(x)) [F(x)]^{n-1}\end{aligned}$$

In general,

$$\begin{aligned}F_k(x) &= \text{prob}(v_k \leq x) \\ &= \text{prob}((\text{number of } v\text{'s that are } \leq x) \geq k) \\ &= \sum_{j=k}^n \frac{n!}{j!(n-j)!} F(x)^j (1 - F(x))^{n-j}\end{aligned}$$

Hence, if F is uniform on $[0, 1]$, i.e., $F(x) = x$, we have

$$\begin{aligned} F_n(x) &= x^n, \quad F_{n-1}(x) = x^n + n(1-x)x^{n-1} \\ f_n(x) &= nx^{n-1}, \quad f_{n-1}(x) = (n-1)n(1-x)x^{n-2}. \end{aligned}$$

Therefore,

$$\begin{aligned} E[v_{(n)}] &= \int_0^1 xnx^{n-1}dx \\ &= \frac{n}{n+1} \end{aligned}$$

and

$$\begin{aligned} E[v_{(n-1)}] &= \int_0^1 x(n-1)n(1-x)x^{n-2}dx \\ &= \frac{n-1}{n+1}. \end{aligned}$$

Now, in a second price auction the expected revenue is the expected second highest value

$$E[R_2] = E[v_{(n-1)}] = \frac{n-1}{n+1},$$

and the revenue in the first price auction is the expected bid of the bidder with the highest value, i.e.,

$$\begin{aligned} E[R_1] &= \frac{n-1}{n} E[v_{(n)}] \\ &= \frac{n-1}{n} \frac{n}{n+1} \\ &= \frac{n-1}{n+1}. \end{aligned}$$

Therefore, both auction forms generate the same expected revenues. This is an illustration of the *revenue equivalence theorem*:

Theorem Any auction with independent private values with a common distribution in which

1. the number of the bidders are the same and the bidders are risk-neutral,
2. the object always goes to the buyer with the highest value,
3. the bidder with the lowest value expects zero surplus,

yields the same expected revenue.

Therefore, all four types of the auctions yield the same expected revenue for the seller in the case of independent private values and risk neutrality. This theorem also allows us to calculate bidding strategies of other auctions. An all-pay auction, for example, satisfies the conditions of the theorem and hence must yield the same expected revenue.

3 Common Values and The Winner's Curse

In a common value auction, bidders have all the same value but each bidder only observes a private signal about the value. Therefore, if a bidder wins the auction, i.e., is the highest bidder, it is likely that the other bidders received worse signals than the winner. In other words, the value of the object conditional on winning is smaller than the unconditional expected value. If this is not taken into account, then the winner might bid an amount more than the actual value of the object, a situation known as the *winner's curse*.

As an example suppose $v = t_1 + t_2$, where v is the common value but bidder i observes only the signal t_i . Assume that each t_i is distributed independently and has a uniform distribution over $[0, 1]$. This, for example, be a takeover battle where the value of the target company is the same but each bidder obtains an independent signal about the value. Suppose that the auction is a first-price sealed bid auction. Denote the strategies by $b_i(t_i)$ and look for an equilibrium in which $b_i(t_i) = at_i$. The expected payoff of player 1 given that player 2 bids according to $b_2(t_2) = at_2$ is given by

$$\begin{aligned} U_1(b_1, t_1) &= E[v - b_1 \mid b_1 > b_2] \text{prob}(b_1 > b_2) \\ &= E[t_1 + t_2 - b_1 \mid b_1 > at_2] \text{prob}(b_1 > at_2) \\ &= E[t_1 + t_2 - b_1 \mid t_2 < b_1/a] \text{prob}(t_2 < b_1/a) \\ &= (t_1 + E[t_2 \mid t_2 < b_1/a] - b_1) \frac{b_1}{a} \\ &= \left(t_1 + \frac{b_1}{2a} - b_1 \right) \frac{b_1}{a} \end{aligned}$$

First order condition is

$$\frac{\partial U_1(b_1, t_1)}{\partial b_1} = \left(t_1 + \frac{b_1}{2a} - b_1 \right) \frac{1}{a} + \frac{b_1}{a} \left(\frac{1}{2a} - 1 \right) = 0,$$

which implies that

$$2 \left(1 - \frac{1}{2a} \right) b_1 = t_1.$$

For $b_1 = at_1$ to be optimal we must have

$$2 \left(1 - \frac{1}{2a} \right) at_1 = t_1,$$

which implies that $a = 1$. Therefore, $b_i(t_i) = t_i$ is a Nash equilibrium.

As a comparison consider the independent private values case where $v_i = t_i + 0.5$. Note that this is the expected value in the above model conditional upon observing t_i (but not conditional upon winning). Let's look for an equilibrium of the form $b_i(t_i) = at_i + c$. The

expected payoff of player 1 to bidding b_1 given that player 2 is using the strategy $at_2 + c$ is

$$\begin{aligned}
U_1(b_1, t_1) &= E[v - b_1] \text{prob}(b_1 > b_2) \\
&= E[t_1 + 0.5 - b_1] \text{prob}(b_1 > at_2 + c) \\
&= (t_1 + 0.5 - b_1) \text{prob}\left(t_2 < \frac{b_1 - c}{a}\right) \\
&= (t_1 + 0.5 - b_1) \frac{b_1 - c}{a}
\end{aligned}$$

if $a \leq b_1 \leq a + c$. Assume that this holds. Then, the first order condition is

$$\frac{\partial U_1(b_1, t_1)}{\partial b_1} = -\frac{b_1 - c}{a} + (t_1 + 0.5 - b_1) \frac{1}{a} = 0$$

which is solved at

$$b_1 = \frac{1}{2}c + \frac{1}{2}t_1 + 0.25.$$

For $b_1(t_1) = at_1 + c$ to be optimal, we must have

$$\frac{1}{2}c + \frac{1}{2}t_1 + 0.25 = at_1 + c,$$

which implies that $a = 0.5$, $c = 0.5$. Therefore,

$$\begin{aligned}
b_1(t_1) &= \frac{t_1}{2} + \frac{1}{2}, \\
b_2(t_2) &= \frac{t_2}{2} + \frac{1}{2}
\end{aligned}$$

constitutes a Bayesian Nash equilibrium (indeed the unique equilibrium) of this auction. (Notice that this satisfies that above restriction $a \leq b_1 \leq a + c$ for all t_1 .) Also note that

$$\frac{t_1}{2} + \frac{1}{2} \geq t_1$$

for all $t_1 \in [0, 1]$, hence there is always underbidding in common value auctions. The reason is that the expected value of the object is smaller conditional upon winning in common value auctions, whereas this value does not depend on the event of winning or not winning.

4 Auction Design

The auctioneer may have different objectives in designing an auction. The government which is privatizing a company, for example, might want to generate the highest revenue from the auction, or might want to make sure that it is efficient, i.e., that the company goes to the bidder with the highest valuation for it, or to a bidder with some other characteristics. Auction theory helps in designing auctions by comparing different auction formats in terms of their equilibrium outcomes. For example, if the objective is to generate the highest revenue, then

different auction formats may be compared on the basis of the expected equilibrium revenues to find the best one. In the case of private, independent values with the same number of risk neutral bidders, revenue equivalence theorem says that the format does not matter, as long as the reserve price is set right. Therefore, the cases where the values are correlated (as in the case of common value auctions), or the bidders are risk averse, auction design becomes a challenging matter. In practice, collusion and entry-deterrence also becomes relevant design problems. Collusion is relevant because revenue equivalence does not hold if there is collusion. Also, remember that the expected revenue from an auction increases in the number of bidders even when the revenue equivalence holds, and hence the auctioneer has an incentive to prevent entry-deterrence.

4.1 Need for a Reserve Price

If there is only one bidder who comes to the auction, the seller will not make any money, unless she sets a reserve price. What is the optimal reserve price? This is similar to the case where the seller is a monopoly and tries to find the optimal price. Assuming that the costs are sunk and therefore the total payoff of the seller is given by the total revenue, the expected payoff is given by

$$\begin{aligned} E[R(p)] &= \text{prob}(\text{sale occurs at price } p) \times p \\ &= \text{prob}(p < v) \times p \\ &= (1 - F(p)) p. \end{aligned}$$

This is maximized when

$$-f(p)p + (1 - F(p)) = 0$$

or

$$p = \frac{1 - F(p)}{f(p)}.$$

So, if F is uniform over $[0,1]$,

$$p = \frac{1 - p}{1}$$

which implies that $p = 1/2$. So, an optimal auction must set a reserve price of 0.5 in this particular case.

4.2 Common Values

We have seen above that first-price sealed bid auction leads to lower bids in the case of common value auctions. In general, if the signals received by the bidders are positively correlated, ascending auction raises more expected revenue than the second-price sealed bid auction, which in turn beats the first-price auction.

4.3 Risk-Averse Bidders

In a second price auction risk aversion does not matter, i.e., the bidders always bid their values. In a first-price auction however, an increase in risk aversion leads to higher bids since it increases the probability of winning at the cost of reducing the value of winning. Therefore, a risk-neutral seller faced with risk-averse bidders prefers the first-price or (descending) Dutch auction to second-price or (ascending) English auctions.

4.4 Practical Concerns³

4.4.1 Collusion

A major concern in practical auction design is the possibility that the bidders explicitly or tacitly collude to avoid higher prices. As an example consider a multi-unit (simultaneous) ascending auction. In such an auction, bidders can use the early stages when prices are still low to signal who should win which objects, and then tacitly agree to stop pushing prices up.

- 1999 Germany spectrum auction: any new bid must exceed the previous one by at least 10 percent. Mannesman bid 18.18 mil. on blocks 1-5 and 20 mil. on blocks 1-6 ($18.18 \times 1.1 \simeq 20$). This was like an offer to T-Mobil (the only other credible bidder) to bid 20 mil. on blocks 1-5 and not bid on blocks 6-10. This is exactly what happened.
- 1996-97 U.S. spectrum auction: U.S. West was competing vigorously with McLeod for lot number 378 - a licence in Rochester, Minnesota. U.S. West bid \$313,378 and \$62,378 for two licences in Iowa in which it had earlier shown no interest, overbidding McLeod who had seemed to be the uncontested high-bidder for these licenses. McLeod got the point that it was being punished for competing in Rochester, and dropped out of that market. Since McLeod made subsequent higher bids on the Iowa licenses, the “punishment” bids cost U.S. West nothing
- A related phenomenon can arise in one special kind of sealed-bid auction, namely a uniform-price auction in which each bidder submits a sealed bid stating what price it would pay for different quantities of a homogenous good, e.g., electricity (that is, it submits a demand function), and then the good is sold at the single price determined by the lowest winning bid. In this format, bidders can submit bids that ensure that any deviation from a (tacit or explicit) collusive agreement is severely punished: each bidder bids very high prices for smaller quantities than its collusively agreed share. Then if any bidder attempts to obtain more than its agreed share (leaving other firms with less than their agreed shares), all bidders will have to pay these very high prices. However,

³This part is based on Paul Klemperer, "What Really Matters in Auction Design," *Journal of Economic Perspectives* 2002, 16(1), 169-189.

if everyone sticks to their agreed shares then these very high prices will never need to be paid. So deviation from the collusive agreement is unprofitable. The electricity regulator in the United Kingdom believes the market in which distribution companies purchase electricity from generating companies has fallen prey to exactly this kind of “implicit collusion.”

Much of the kind of behavior discussed so far is hard to challenge legally. Indeed, trying to outlaw it all would require cumbersome rules that restrict bidders’ flexibility and might generate inefficiencies, without being fully effective. It would be much better to solve these problems with better auction designs.

4.4.2 Entry Deterrence

The second major area of concern of practical auction design is to attract bidders, since an auction with too few bidders risks being unprofitable for the auctioneer and potentially inefficient. Ascending auctions are often particularly poor in this respect, since they can allow some bidders to deter the entry, or depress the bidding, of rivals. In an ascending auction, there is a strong presumption that the firm which values winning the most will be the eventual winner, because even if it is outbid at an early stage, it can eventually top any opposition. As a result, other firms have little incentive to enter the bidding, and may not do so if they have even modest costs of bidding.

- Glaxo’s 1995 takeover of the Wellcome drugs company. After Glaxo’s first bid of 9 billion pounds, Zeneca expressed willingness to offer about 10 billion pounds if it could be sure of winning, while Roche considered an offer of 11 billion pounds. But certain synergies made Wellcome worth a little more to Glaxo than to the other firms, and the costs of bidding were tens of millions of pounds. Eventually, neither Roche nor Zeneca actually entered the bidding, and Wellcome was sold at the original bid of 9 billion pounds, literally a billion or two less than its shareholders might have received. Wellcome’s own chief executive admitted “...there was money left on the table”.

4.4.3 Solutions

Much of our discussion has emphasized the vulnerability of ascending auctions to collusion and predatory behavior. However, ascending auctions have several virtues, as well.

- An ascending auction is particularly likely to allocate the prizes to the bidders who value them the most, since a bidder with a higher value always has the opportunity to rebid to top a lower-value bidder who may initially have bid more aggressively.
- If there are complementarities between the objects for sale, a multi-unit ascending auction makes it more likely that bidders will win efficient bundles than in a pure sealed-bid auction in which they can learn nothing about their opponents’ intentions.

- Allowing bidders to learn about others' valuations during the auction can also make the bidders more comfortable with their own assessments and less cautious, and often raises the auctioneer's revenues if information is correlated.

A number of methods to make the ascending auction more robust are clear enough. For example, bidders can be forced to bid "round" numbers, the exact increments can be prespecified, and bids can be made anonymous. These steps make it harder to use bids to signal other buyers. Lots can be aggregated into larger packages to make it harder for bidders to divide the spoils, and keeping secret the number of bidders remaining in the auction also makes collusion harder. But while these measures can be useful, they do not eliminate the risks of collusion or of too few bidders. An alternative is to choose a different type of auction.

In a standard sealed-bid auction (or "first-price" sealed-bid auction), each bidder simultaneously makes a single "best and final" offer, so collusion is much harder than in an ascending auction because firms are unable to retaliate against bidders who fail to cooperate with them. Tacit collusion is particularly difficult since firms are unable to use the bidding to signal.

From the perspective of encouraging more entry, the merit of a sealed-bid auction is that the outcome is much less certain than in an ascending auction. An advantaged bidder will probably win a sealed-bid auction, but it must make its single final offer in the face of uncertainty about its rivals' bids, and because it wants to get a bargain its sealed-bid will not be the maximum it could be pushed to in an ascending auction. So "weaker" bidders have at least some chance of victory, even when they would surely lose an ascending auction. It follows that potential entrants are likely to be more willing to enter a sealed-bid auction than an ascending auction.

A solution to the dilemma of choosing between the ascending (often called "English") and sealed-bid (or "Dutch") forms is to combine the two in a hybrid, the "Anglo-Dutch", which often captures the best features of both, and was first described and proposed in Klemperer (1998. "Auctions with Almost Common Values." *European Economic Review*. 42, pp. 757-69.).

In an Anglo-Dutch auction the auctioneer begins by running an ascending auction in which price is raised continuously until all but two bidders have dropped out. The two remaining bidders are then each required to make a final sealed-bid offer that is not lower than the current asking price, and the winner pays his bid.

Good auction design is not "one size fits all" and must be sensitive to the details of the context.

Extensive Form Games with Perfect Information

Levent Koçkesen

1 Extensive Form Games

So far we have assumed that players, when taking their actions, either did so simultaneously, or without knowing the action choice of the other players. Although, this modelling assumption might be appropriate in some settings, there are many situations in the world of business and politics that involve players moving sequentially after observing what the other players have done. For example, a bargaining situation between a seller and a buyer may involve the buyer making an offer and the seller, after observing the buyer's offer, either accepting or rejecting it. Or imagine an incumbent senator deciding whether to run an expensive ad campaign for the upcoming elections and a potential challenger deciding whether to enter the race or not, after observing the campaign decision of the incumbent. Both of these situations involve a player choosing an action after observing the action of the other player.

The **extensive form** of a game, as opposed to the strategic form, provides a more appropriate framework to analyze certain interesting questions that arise in strategic interactions that involve sequential moves.

1.1 Game Trees

As you now very well know, strategic form of a game has three ingredients: (1) the set of players, (2) the set of actions, and (3) the payoff functions. The extensive form provides a richer specification of a strategic interaction by specifying **who** moves **when** doing **what** and with **what information**. The easiest way to represent an extensive form game is to use a **game tree**, which is a multi-person generalization of a decision tree.

To illustrate, let us go back to the bargaining example above and assume that the buyer moves first by offering either \$500 or \$100 for a product that she values \$600. The seller, for whom the value of the object is \$50, responds by either accepting (A) or rejecting (R) the offer. We can represent this situation by the game tree in Figure 1.

Game trees are made up of

- nodes
- branches
- information sets




Figure 1: Bargaining Game

- player labels
 - action labels
- and
- payoffs.

Nodes are of two types: **Decision nodes** represent the points in the game at which players make a decision, i.e., choose an action, or a strategy in general. As any other tree, a game tree has a root and it is useful to distinguish the root, which we will call the **initial node**, from the other decision nodes (it is represented by an open circle whereas all the other nodes are represented by closed circles). To each decision node, including the initial node, one, and only one, player label is attached, to indicate who moves at that particular decision node. The second type of nodes are called **terminal nodes** and at these nodes the game is over and nobody takes any action anymore. To each terminal node a **payoff** vector is appended.

From each decision node, one or more branches emanate, each branch representing an action that can be taken by the player who is to move at that node. Each such branch is labelled with the **action** that it represents. A branch either leads to another decision node or to a terminal node.

The last component that we have to talk about is the **information sets**. Information sets tell us what the players know when they are making a decision. They are collections of decision nodes of a player that cannot be distinguished from the perspective of that player. We can illustrate it using the bargaining example under the assumption that the seller, somehow, does not observe the buyer's offer before deciding whether to accept or reject it. We depict this informational assumption by connecting the two decision nodes of the seller with a dashed line (see Figure 2).

Notice that the actions available to the seller at the two nodes that are in the same information set must be the same, otherwise the seller would be able to distinguish between them by just looking at the actions available to her.




Figure 2: Bargaining Game with Imperfect Information

In this section we will deal with extensive form games with perfect information in which every player can distinguish between any two decision nodes and hence we will not have to worry about information sets.

1.2 Strategies

Strategies in a strategic form game are either action choices or probability distributions over actions. In an extensive form game, description of a strategy is more involved since players may have to choose actions at several points in the game. Therefore, a **pure strategy** of a player in an extensive form game has to specify an action choice at every decision node of that player. In that sense, a strategy is a complete plan of action, so complete that if it was handed over to a computer, the computer would know what to do under every contingency. We denote a pure strategy of player i by s_i , and the set of all pure strategies by S_i .

For example, in the extensive form game in Figure 1, a pure strategy for the buyer is easy enough: it has to specify what price to offer at the initial node. A pure strategy for the seller, on the other hand, has to specify an action at each decision node she may be called upon to move. So, the buyer has two pure strategies available to her: 100 and 500, and hence $S_B = \{100, 500\}$. The seller, however, has four pure strategies: (1) AA , (2) AR , (3) RA , (4) RR , and hence $S_S = \{AA, AR, RA, RR\}$.

The extensive form strategies sometimes lead to confusion. Let us try to illustrate why, by looking at the extensive form game in Figure 3.




Figure 3: Another Extensive Form Game

A strategy for player 1 in this game has to specify an action at every decision node she has, and there are two such nodes. She, therefore, has four strategies: LL' , LR' , RL' , RR' . Notice that the first two strategies specify an action even at player 1's second decision node which would not be reached if those strategies were implemented. The reason why, will become clear in the next section, after we analyze the optimal behavior of players. For now, let us look at the game tree of the senate-race game (see Figure 4) to further illustrate the concepts introduced so far.




Figure 4: Senate-Race Game

In this game $S_I = \{A, N\}$ and $S_C = \{ee, en, ne, nn\}$.

2 Backward Induction Equilibrium

As in the strategic form games, the equilibrium concept in extensive form games is based upon the idea that each player plays a best response to the play of the other players. The difference is that we now require strategies to be optimal at every step in the game. The **backward induction equilibrium** is an algorithm that results in a recommendation of an action choice at every decision node with the property that if every player follows those recommendations their strategies would be optimal at every decision node they *may* be called upon to move. This will also result in a path of play (i.e., a sequence of branches) which will be called the **backward induction outcome**.

The algorithm is really simple. You, the game theorist, go to the final decision nodes and determine the best action available to the players who are to move at those nodes. Since there is no more moves after players make their moves at these decision nodes, this boils down to choosing the action that leads to the highest payoff for the player who is moving. (If there is a tie between two actions that lead to the highest payoff, you may simply choose one of them.) After you have done that, you prune all the actions that are not chosen (or just indicate the ones that are chosen by an arrow-head) and go to the penultimate decision nodes to determine the optimal action at those nodes. You continue in this manner until you reach to the initial node and determine the optimal action there.

For example, in the bargaining game we start with the seller's decision nodes which are the final decision nodes in the game tree. Since accepting both offers is optimal we mark the

branches labelled A by arrow-heads. Once we do that, it is easily seen that the best action for player 1 is to offer \$100. Therefore, the backward induction equilibrium of the bargaining game is $(100, AA)$ and the backward induction outcome is $(100, A)$. (See Figure 5) The backward induction equilibrium of the senate race game is (A, ne) and its backward induction outcome is (A, n) . (See Figure 6).




Figure 5: Bargaining Game




Figure 6: Senate-Race Game

As an exercise verify that the backward induction equilibrium of the game in Figure 3 is (LR', r) . This example illustrates why player 1's strategy had to specify an action even after she has previously chosen L . Whether L is optimal or not for player 1 depends on what she believes that player 2 will do. If she believes that player 2 is going to choose l , then L is not optimal. But, whether player 2 will choose l or not depends on what player 2 believes that player 1 is going to do in her last decision node. Therefore, to determine the optimal action for player 1 at her first decision node, we have to specify what she intends to do at her last decision node.




Figure 7: Another Extensive Form Game

2.1 Commitment and Mover Advantages

The bargaining and the senate-race games illustrate an important phenomenon that arise in many extensive form games, i.e., **the power of commitment**. Suppose that the seller could, somehow, commit herself to accepting only the offer 500 and that this is known to the buyer. Now, given that knowledge, the best that the buyer can do is to actually offer 500, because otherwise her offer will be rejected and she will receive 0, whereas offering 500 gives her 100. Therefore, public, and credible, commitments could increase a player's payoff in an extensive form game. Notice that this is similar to eliminating action A after the offer 100. This is in stark contrast to the single individual decision making problems where eliminating an action can never improve one's payoff.

Similarly, in the senate-race game, if the challenger could publicly commit to entering the race irrespective of the campaign decision of the incumbent, the best thing the incumbent could do would be not to run campaign ads and hence the challenger would respond by entering the race and obtaining a payoff of 4 rather than 3 that she was getting in the backward induction equilibrium.¹

Another interesting phenomenon that arise in certain extensive form games is that of **first mover advantage**. For example, in the senate-race game, when the incumbent moves first, both players obtain a payoff of 3 in the backward induction equilibrium. Now, let us change the order of the moves so that it is the challenger who moves first so that we obtain the game tree depicted in Figure 8.




Figure 8: Modified Senate-Race Game

The backward induction equilibrium of this game is (e, NN) which yields a payoff of 4 to the challenger and a payoff of 2 to the incumbent. Therefore, if they had the chance, both players would prefer to move first in this game. This is similar to the idea behind the power of commitment. By choosing e the challenger commits herself to entering whatever the incumbent does.

However, not all games have a first mover advantage. Quite to the contrary, some games have **second mover advantage**. Consider a game in which the incumbent (who belongs

¹See Thomas Schelling (1960), *The Strategy of Conflict*, for an excellent account of the idea of credible commitments.

to a rightist party) and the challenger (who belongs to a leftist party) in a senate race are choosing political platforms; either a leftist or a rightist one. Suppose that if both of them choose the same platform the incumbent wins the elections, whereas if they choose different platforms it is the challenger who wins. The candidates mostly care about winning, but they also would like to win (or lose) without compromising their political views. The game tree in Figure 9 depicts the situation if it is the incumbent who moves first, whereas the one in Figure 10 reverses the order of moves. Verify that this game exhibits second mover advantage.




Figure 9: Senate-Race Game II




Figure 10: Modified Senate-Race Game II

3 Game Trees: A More Formal Treatment

A game tree is a collection of nodes, called T , and a binary relation between the nodes called a *precedence* relation, denoted $\succ$. Given two nodes α and β in the game tree, $\alpha \succ \beta$ means that α precedes β . Using this relation, we can define the set of predecessors of α as

$$P(\alpha) = \{t \in T : t \succ \alpha\}$$

and the set of successors as

$$S(\alpha) = \{t \in T : \alpha \succ t\}.$$

The set $P(\alpha)$ is simply the set of nodes from which one can go (through a sequence of branches) to α . Similarly, the set of successors of α is the set of nodes to which one can go starting from α .

The precedence relation $\succ$

1. is asymmetric, i.e., there exists no $\alpha, \beta \in T$ such that $\alpha \succ \beta$ and $\beta \succ \alpha$;
2. is transitive, i.e., $\alpha \succ \beta$ and $\beta \succ \gamma$ implies $\alpha \succ \gamma$;
3. there is a common predecessor to any two non-initial nodes, i.e., for all $\alpha, \beta \in T$, with $P(\alpha) \neq \emptyset$ and $P(\beta) \neq \emptyset$, there exists a $\gamma \in T$ such that $\gamma \in P(\alpha)$ and $\gamma \in P(\beta)$.
4. and satisfies the following property

If $\alpha \succ \gamma$ and $\beta \succ \gamma$, then either $\alpha \succ \beta$ or $\beta \succ \alpha$.

The first two conditions guarantee that there are no cycles in the game tree, while the third condition guarantees that there is a unique initial node. The last condition guarantees that starting from any node there is a unique path back to the initial node.

Theorem 1 *Kuhn's (Zermelo's Theorem). Every finite extensive form game with perfect information has a backward induction equilibrium.*

Proof. Omitted.

4 Strategic Form of an Extensive Form Game

The strategic form is given by

1. The set of players N ,
and for each player i
2. The set of strategies S_i ,
3. The payoff function,

$$u_i : S \rightarrow \mathbf{R}$$

where $S = \times_{i \in N} S_i$ is the set of all strategy profiles.

So, the only difference from the standard definition of a strategic form game is the use of strategies rather than actions.

As an illustration, let us find the strategic form of the bargaining game. The set of players is $N = \{B, S\}$, the set of strategies are $S_B = \{100, 500\}$ and $S_S = \{AA, AR, RA, RR\}$. The payoff functions are represented in the following bimatrix

	AA	AR	RA	RR
100	500, 50	500, 50	0, 0	0, 0
500	100, 450	0, 0	100, 450	0, 0

Definition 2 A strategy profile $s^* \in S$ is a Nash equilibrium if for each player i

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \text{ for all } s_i \in S_i$$

or equivalently, if for each player i

$$s_i^* \in B_i(s_{-i}^*).$$

Therefore, the above bargaining game has three Nash equilibria $(100, AA)$, $(100, AR)$, and $(500, RA)$. Notice that the first two Nash equilibria result in the same outcome as does the backward induction equilibrium, i.e., $(100, A)$, whereas the third one results in the outcome $(500, A)$. This last equilibrium, however, is sustained by an **incredible threat** by the seller, i.e., the threat that she will not accept the offer of \$100. This threat is not credible because, if it was tested by the buyer, i.e., the buyer were to offer \$100, then the seller would actually find it in her interest to accept the offer.

Backward induction equilibrium concept eliminates equilibria based upon incredible threats by demanding players to be rational at every point in the game, a property that we call **sequential rationality**. Sequential rationality is stronger than just requiring the strategies to be best responses to the strategies of the other players, i.e., stronger than the rationality requirement behind the Nash equilibrium concept. For example, in the bargaining game above, the strategy RA is a best response to the offer of \$500, but is not sequentially rational, because it specifies the seller to reject the offer of \$100, and this is not rational at the decision node of the seller following the offer of \$100.

5 Extensive Form Games with Imperfect Information

In the previous section we have analyzed extensive form games with perfect information where every player had a perfect knowledge of what had happened previously in the game, i.e., each player observed the previous moves made by the other players. In this section we will relax this assumption and allow the possibility that some of the previous moves by other players are not observed when a player is called upon to move. Such games are called **extensive form games with imperfect information**.

In extensive form games with imperfect information, the notion of information sets, which we have introduced in the last section becomes crucial. An **information set** of player i is a collection of decision nodes of player i that cannot be distinguished by player i . Therefore, if the game reaches to any of the nodes in an information set of a player, that player does not know which of the nodes in that information node has actually been reached.

As an example consider the bargaining game with imperfect information (see Figure 2). In this game there is one information set of the seller that contains the decision nodes following

the offers 100 and 500. When the seller is called upon to move, she does not know which of the two offers have been made, i.e., which of the two decision nodes in the information set has been reached. The strategy sets are given by $S_B = \{100, 500\}$ and $S_S = \{A, R\}$ and hence we have the following strategic form of this game

	A	R
100	500, 50	0, 0
500	100, 450	0, 0

The unique Nash equilibrium of this game is therefore $(100, A)$, the same outcome as the backward induction equilibrium outcome of the bargaining game with perfect information! The reason why we have a unique Nash equilibrium outcome in this game is that we have eliminated the seller's ability of making a non-credible threat of rejecting the offer of \$100.

We may think of extensive form with imperfect information as a generalization of extensive form with perfect information. In the latter, all the information sets are singletons, i.e., they each contain a single node, whereas in the former there is at least one information set that contains more than one node.

As an another example consider the following entry-game. Suppose Pepsi is currently the sole provider in a market, say in Bulgaria. Coke is considering to enter the market. If Coke enters, both firms simultaneously decide whether to act tough (T) or accommodate (A). This leads to an extensive form game with imperfect information whose game tree representation is given in Figure 11, where the first number in a payoff vector belongs to Coke and the second to Pepsi.




Figure 11: Entry-game.

In this game $S_C = \{OT, OA, ET, EA\}$ and $S_P = \{T, A\}$, and hence we have the following

strategic form:

	T	A
OT	0, 5	0, 5
OA	0, 5	0, 5
ET	-2, -1	0, -3
EA	-3, 1	1, 2

There are three Nash equilibria of this game: (OT, T) , (OA, T) , (EA, A) . In the second Nash equilibrium Coke is supposed to accommodate and Pepsi is supposed to act tough, following Coke entering the market. Is that reasonable? In other words, suppose, the game actually reached that stage, that is Coke actually entered. Now, is (A, T) a reasonable outcome? One way of asking the same question is to check if both players are acting rationally, i.e., best responding to each other's strategies, conditional upon Coke entering the market. Notice that conditional upon Coke entering the market we have the following "game"




Figure 12: Entry-"game".

which has the following strategic form

		Pepsi	
		T	A
Coke	T	-2, -1	0, -3
	A	-3, 1	1, 2

If Coke anticipates Pepsi to play T , then its best response is T as well, not A . (Neither is T a best response for Pepsi to A .) Therefore, to the extent that we regard only Nash equilibrium outcomes as reasonable, we conclude that (A, T) is not reasonable. In contrast, the post-entry behavior of both players are rational in equilibria (OT, T) and (EA, A) .

5.1 Subgames and Subgame Perfect Equilibrium

Subgame perfect equilibrium is a generalization of the backward induction equilibrium to extensive form games with imperfect information. To define subgame perfect equilibrium we have to first define a subgame.

Definition 3 A *subgame* is a part of the game tree such that

1. it starts at a single decision node,
2. it contains every successor to this node,
3. if it contains a node in an information set, then it contains all the nodes in that information set.

It is conventional to treat the entire game as a subgame and call all the other subgames **proper subgames**. For example, the entry-game given in figure 11 has two subgames: the game itself and the subgame which starts after Coke enters the market. Of course, only the latter is a proper subgame.

Given a subgame g , let us denote the restriction of a strategy s_i to that subgame g by $s_i|_g$. For example, if we denote the post-entry subgame in the entry-game by e (this subgame is given in figure 12), then $OT|_e = T$, $EA|_e = A$, etc.

Definition 4 A strategy profile s^* in an extensive form game Γ is a **subgame perfect equilibrium** (SPE) if for every subgame g of Γ , $s^*|_g$ is a Nash equilibrium of g .

Therefore, there are two SPE of the entry-game: (OT, T) and (EA, A) .

We can now obtain a better insight into the difference between subgame perfect equilibrium (or backward induction equilibrium) and Nash equilibrium by using the language of subgames. We first have to distinguish between subgames that can be reached by a strategy profile and those that cannot be reached. A subgame can be **reached** under the strategy profile $s \in S$ if, when the strategy profile is implemented, the initial node of the subgame will actually be reached. Otherwise, we say that the subgame cannot be reached under the strategy profile s . A strategy profile s^* is a Nash equilibrium if every player plays a best response to the strategies of the other players in every subgame that can be reached under s^* . In contrast, a strategy profile s^* is a SPE if every player plays a best response to the strategies of the other players in every subgame, i.e., even in those subgames that cannot be reached under s^* . In other words, Nash equilibrium demands rationality in only those subgames that can be reached in equilibrium, whereas SPE demands rationality in every subgame, and this latter form of rationality is called **sequential rationality**.

As an exercise consider the game in figure 3 and find its Nash equilibria and SPE. Verify that there are Nash equilibria in which one of the players do not behave sequentially rationally, whereas in all SPE both players act sequentially rationally.

Extensive Form Games: Applications

Levent Koçkesen

1 Bargaining

1.1 Ultimatum Bargaining

Two players, A and B, bargain over a cake of size 1. Player A makes an offer $x_A \in [0, 1]$ to player B. If player B accepts the offer (Y), agreement is reached and player A receives x_A and player B receives $1 - x_A$. If player B rejects the offer (N), both players receive a payoff of zero. This can be modelled as an extensive form game with perfect information. However, it is not a finite game as A has infinitely many actions.

We can use backward induction to find the subgame perfect equilibrium (SPE) of this game. Consider a subgame that follows A's offer of x . If $x < 1$, then B's optimal action is to accept the offer. If, on the other hand, $x = 1$, then both accepting and rejecting are optimal. First, suppose that B accepts any offer $x \in [0, 1]$. In this case, clearly, the optimal offer by A is $x_A^* = 1$. So, one SPE is $(1, s_2^*(x))$ where

$$s_2^*(x) = Y \text{ for all } x \in [0, 1].$$

Now suppose B accepts only offers that are strictly smaller than one. What is A's optimal offer in this case? Could offering 1 be optimal? No, because this will be rejected by B resulting a payoff of zero for A. Player A could deviate and offer something smaller than one and obtain a positive payoff instead. Could offering something strictly smaller than 1 be optimal? No! To see why, suppose $x < 1$ is an optimal offer. This will be accepted by B and give player A a payoff of x . However, player A can deviate and offer $x + \varepsilon$, with $0 < \varepsilon < 1 - x$ and hence $x + \varepsilon < 1$, which will be accepted by B and give player A a payoff of $x + \varepsilon$ which is strictly greater than x . Therefore, the unique SPE is $(1, s_2^*(x))$ where

$$s_2^*(x) = Y \text{ for all } x \in [0, 1]$$

and the unique SPE outcome is $(1, Y)$.

1.2 Alternating Offers Bargaining

1.2.1 Preliminaries

Two players, A and B, bargain over a cake of size 1. At time 0 player A makes an offer $x_A \in [0, 1]$ to player B. If player B accepts the offer, agreement is reached and player A

receives x_A and player B receives $1 - x_A$. If player B rejects the offer, she makes a counteroffer $x_B \in [0, 1]$ at time 1. If this counteroffer is accepted by player A, then player B receives x_B and player A receives $1 - x_B$. Otherwise, player A makes another offer at time 2. This process continues indefinitely until a player accepts an offer.

If the players reach an agreement at time t on a partition that gives player i a share x_i of the cake, then player i 's payoff is $\delta_i^t x_i$, where $\delta_i \in (0, 1)$ is player i 's discount factor. If players never reach an agreement, then each player's payoff is zero.

1.2.2 Stationary No-delay Equilibrium

We will first characterize the SPE with the following two properties, and then show that all SPE have these properties.

1. **(No Delay)** All equilibrium offers are accepted.
2. **(Stationarity)** A player makes always the same offer in equilibrium.

Let x_i^* denote the equilibrium offer by player i . Given properties 1 and 2, the current present value of rejecting an offer x_A^* is $\delta_B x_B^*$ for player B. This implies that in equilibrium

$$1 - x_A^* = \delta_B x_B^*. \quad (1)$$

Similarly

$$1 - x_B^* = \delta_A x_A^*. \quad (2)$$

Therefore, there is a unique solution

$$x_A^* = \frac{1 - \delta_B}{1 - \delta_A \delta_B} \quad (3)$$

$$x_B^* = \frac{1 - \delta_A}{1 - \delta_A \delta_B} \quad (4)$$

Thus, there exists at most one SPE satisfying the two properties. But we still have to verify there exists such an equilibrium. Consider the strategy profile (s_A^*, s_B^*) defined as

Player A: Always offer x_A^* , accept any x_B with $1 - x_B \geq \delta_A x_A^*$

Player B: Always offer x_B^* , accept any x_A with $1 - x_A \geq \delta_B x_B^*$.

Before we prove that this strategy profile is a SPE we state the following proposition

Proposition 1 (One-Deviation Property). *Let Γ be a finite horizon extensive form game with perfect information. The strategy profile s^* is a SPE of Γ if and only if for every player $i \in N$ and for every subgame g of Γ , in which player i moves at the initial node of g there exists no profitable deviation by player i which differs from s_i^* only in the action specified at the initial node of g .*

Remark 1 *It is possible to show that if the payoffs of an **infinite horizon** game satisfies a certain regularity condition (continuity at infinity: see Fudenberg and Tirole, 1991, p. 110), then the one-deviation property holds for infinite horizon games as well.]*

Proposition 2 *One-deviation property holds for the Rubinstein bargaining game.*

Proposition 3 *(s_A^*, s_B^*) is a SPE of the alternating offers bargaining game.*

Proof. Consider any period when A has to make an offer. Her payoff to s_A^* is x_A^* . If A offers $x_A < x_A^*$ then

$$1 - x_A > \delta_B x_B^*$$

by equation (1) and hence B accepts any such offer which gives A a payoff less than x_A^* . If she offers $x_A > x_A^*$, then B rejects and offers x_B^* , A accepts giving her a payoff of

$$\delta_A (1 - x_B^*) < x_A^*$$

by equation (1). Therefore, there is no profitable one-shot deviation in any subgame starting with her offer.

Now, consider subgames starting with player A responding. If player A rejects offer x_B with $1 - x_B \geq \delta_A x_A^*$, then she will offer x_A^* herself and get $\delta_A x_A^*$. So this is not a profitable deviation.

By a symmetric argument, it follows that player B's strategy is optimal in every subgame as well. ■

1.3 Unique Subgame Perfect Equilibrium

Theorem 4 *The strategy profile s^* is the unique SPE.*

Proof. Let Γ_i denote any subgame that starts with player i making an offer. Clearly, all such subgames are strategically equivalent (since preferences are stationary, i.e., does not depend on calendar time) and thus all have the same SPE. Let G_i denote the set of SPE payoffs of player i in any subgame Γ_i and let¹

$$\begin{aligned} M_i &= \max G_i, \\ m_i &= \min G_i. \end{aligned}$$

Lemma 5 *There exists a unique SPE payoff profile of Γ_A given by $(x_A^*, 1 - x_A^*)$ and a unique SPE payoff profile of Γ_B given by $(x_B^*, 1 - x_B^*)$.*

¹It is possible that $\max G_i$ and $\min G_i$ do not exist. For example if $G_i = (0, 1)$, $\max G_i$ and $\min G_i$ do not exist. However, the theorem is still true with max and min replaced with sup (supremum) and inf (infimum).

Proof of Lemma 5.

Claim 1. $m_i \geq 1 - \delta_j M_j$, $i \neq j$.

Proof of Claim 1. First note that player B accepts any offer x_A such that $1 - x_A > \delta_B M_B$. So, if there exists an equilibrium of Γ_A yielding $u_A < 1 - \delta_B M_B$, player A can profitably deviate from such an equilibrium by offering x_A such that $u_A < x_A < 1 - \delta_B M_B$. $\parallel$

Claim 2. $M_i \leq 1 - \delta_j m_j$, $i \neq j$.

Proof of Claim 2. Player B rejects any offer which gives her less than $\delta_B m_B$ and following rejection she never offers more than $\delta_A M_A$. Therefore,

$$M_A \leq \max \left\{ \underbrace{1 - \delta_B m_B}_{\text{max when B accepts}}, \underbrace{\delta_A^2 M_A}_{\text{max when B rejects}} \right\}$$

and hence

$$M_A \leq 1 - \delta_B m_B.$$

$\parallel$

Claims 1 and 2 imply that

$$m_A \geq 1 - \delta_B M_B \tag{5}$$

$$m_B \geq 1 - \delta_A M_A \tag{6}$$

$$M_A \leq 1 - \delta_B m_B \tag{7}$$

$$M_B \leq 1 - \delta_A m_A \tag{8}$$

From 6 we get

$$-\delta_B m_B \leq -\delta_B (1 - \delta_A M_A)$$

and

$$1 - \delta_B m_B \leq 1 - \delta_B (1 - \delta_A M_A)$$

which together with 7 implies that

$$M_A \leq 1 - \delta_B (1 - \delta_A M_A)$$

or

$$M_A \leq \frac{1 - \delta_B}{1 - \delta_A \delta_B}. \tag{9}$$

By 8 we have

$$1 - \delta_B M_B \geq 1 - \delta_B (1 - \delta_A m_A)$$

which together with 5 implies that

$$m_A \geq 1 - \delta_B (1 - \delta_A m_A)$$

or

$$m_A \geq \frac{1 - \delta_B}{1 - \delta_A \delta_B}. \quad (10)$$

Since $M_A \geq m_A$, 9 and 10 imply that

$$M_A = m_A = \frac{1 - \delta_B}{1 - \delta_A \delta_B}. \quad (11)$$

Similarly,

$$M_B = m_B = \frac{1 - \delta_A}{1 - \delta_A \delta_B}. \quad (12)$$

Therefore, the unique payoff profile in Γ_A is $(x_A^*, 1 - x_A^*)$ and the unique payoff profile in Γ_B is $(x_B^*, 1 - x_B^*)$.||

We can now complete the proof of the theorem by using Lemma 5. We first show that all equilibrium offers are accepted in any SPE. Suppose there exists a SPE in which player A's offer is rejected. By Lemma 5, A's equilibrium payoff in this subgame is x_A^* . But by Lemma 5, A's payoff in subgame following rejection is $(1 - x_B^*)$, and hence, the equilibrium payoff of A in the subgame in which her offer is rejected must be $\delta_A (1 - x_B^*)$. But, this implies

$$x_A^* = \delta_A (1 - x_B^*) = \delta_A^2 x_A^*,$$

a contradiction. Similarly, player B's equilibrium offers must be accepted.

Second, we show that in all SPE A offers x_A^* and B offers x_B^* . Suppose A offers $x_A > x_A^*$ in a SPE of a Γ_A . This offer must be rejected by B in equilibrium, because otherwise B would get less than $1 - x_A^*$ in that subgame which contradicts Lemma 5. This, in turn, contradicts that no equilibrium offer is rejected. Now suppose $x_A < x_A^*$ in a SPE of a Γ_A . This offer, too, must be rejected by B, because otherwise A would get less than x_A^* in that subgame, contradicting Lemma 5. So, A must be offering x_A^* in all SPE. Similarly, B must always be offering x_B^* .

Since there is a unique SPE satisfying these properties, as proved in Proposition 3, the proof is complete.■

1.4 Properties of the Equilibrium

(1) Equilibrium is Unique and Efficient

This is the case for $\delta_i < 1$. That is, there has to be some friction. Otherwise there exists a continuum of equilibria, including inefficient ones.

(2) Bargaining Power

Note that the share of player A in the unique SPE is

$$\pi_A = x_A^* = \frac{1 - \delta_B}{1 - \delta_A \delta_B}$$

and that of B is

$$\pi_B = 1 - x_A^* = \frac{\delta_B(1 - \delta_A)}{1 - \delta_A\delta_B}$$

and hence the share of player i is increasing in δ_i and decreasing in δ_j . The bargaining power comes from patience. The more patient a player is, the higher her share.

If the payers are equally patient, i.e., $\delta_A = \delta_B = \delta$, then

$$\pi_A = \frac{1}{1 + \delta} > \frac{\delta}{1 + \delta} = \pi_B$$

In other words, there is a first mover advantage.

The first mover advantage disappears as $\delta \rightarrow 1$.

$$\begin{aligned} \lim_{\delta \rightarrow 1} \pi_A &= \frac{1}{2} \\ \lim_{\delta \rightarrow 1} \pi_B &= \frac{1}{2}. \end{aligned}$$

Infinitely Repeated Games

Levent Koçkesen

1 Motivation

Many interactions in the real world have an ongoing structure and in many such situations people consider their long-term payoffs in addition to the short-term gains. This might lead people to behave in ways different from how they would if the interactions were one-shot rather than long-term. Consider the following prisoners' dilemma game.

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

Remember that in this game defecting (D) for both players is the unique Nash equilibrium (and also the strictly dominant strategy equilibrium). So, if this game is played only once, game theory strongly suggests that the outcome will be (D, D) , which is suboptimal, since the cooperative outcome (C, C) gives both players a strictly higher payoff. However, if this game is played repeatedly between two players, then they may be inclined to cooperate, rather than defect, if they think they will be punished in the future for defecting.

Theory of repeated games analyzes the types of outcomes, behavior, and norms that can be supported as Nash equilibrium and subgame perfect equilibrium outcomes in repeated interactions. Rather than presenting this large body of literature, we will present some examples and indicate how they generalize to other repeated interactions.

2 Preliminaries

Let $G = (N, (A_i), (u_i))$ be an n -player finite strategic form game. We will call G the **stage game**. For example, G might be the prisoners' dilemma (PD) game given above.

An **infinitely repeated game** is defined by the following elements. The stage game is played at each discrete time period $t = 1, 2, \dots$, and at the end of each period the action choice of each player is revealed to everybody. A **history** in time period t is simply a sequence of action profiles from period 1 through period $t - 1$, i.e.,

$$h^t = (a^0, a^1, a^2, \dots, a^{t-1}), \text{ for } t = 1, 2, \dots$$

where we take a^0 to be the empty history (i.e., nothing has happened so far). For example, in the PD game a possible fifth period history is $(a^0, (C, C), (C, C), (D, C), (D, D))$. [We will usually omit the empty history in this specification as a convention, and write $((C, C), (C, C), (D, C), (D, D))$.] The set of period t histories is then given by

$$H^t = A^{t-1}, \text{ for } t = 1, 2, \dots$$

where we again set $A^0 = a^0$. The set of all second period histories in PD game is

$$\begin{aligned} H^2 &= A \\ &= \{(C, C), (C, D), (D, C), (D, D)\} \end{aligned}$$

and the set of all period three histories is

$$\begin{aligned} H^3 &= A^2 = A \times A \\ &= \{(C, C), (C, D), (D, C), (D, D)\} \times \{(C, C), (C, D), (D, C), (D, D)\} \end{aligned}$$

etc.. A history is terminal if and only if it is infinite. In other words a terminal history is in the form of $(a^0, a^1, a^2, \dots)$. Notice that each nonterminal history starts a subgame in the repeated game.

After any nonterminal history each player $i \in N$ simultaneously chooses an action in A_i . Therefore, a **pure strategy** s_i of player i is a sequence of functions that assign an action in A_i to every history h^t ; $s_i(h^t)$ denotes the action choice of player i after history h^t . Therefore, a strategy for player i is given by

$$s_i = (s_i(a^0), s_i(a^0, a^1), \dots, s_i(a^0, a^1, \dots, a^{t-1}), \dots)$$

For example, in the PD game a strategy may specify

$$\begin{aligned} s_i(a^0) &= C \text{ and} \\ s_i(a^0, a^1, \dots, a^{t-1}) &= \begin{cases} C, & \text{if } a_j^\tau = C, j \neq i, \text{ for } \tau = 1, 2, \dots, t-1 \\ D, & \text{otherwise} \end{cases} \end{aligned}$$

This strategy instructs player i to start with playing C and continuing doing so unless the opponent has played D in the past, in which case, player i plays D forever. (This strategy is also called **grim-trigger** strategy, because defection is triggered by the opponent's defection and because punishment is unrelenting). We denote the set of all pure strategies for player i by S_i . The set of all strategy profiles are denoted S . A strategy profile $s = (s_1, \dots, s_n)$ induces a terminal history in the obvious manner. For example, if both players adopt the grim-trigger strategy defined above the outcome will be cooperation every period.

The last thing that we have to define is **payoff functions**. Since, the only terminal histories are the infinite histories and each period's payoff is the payoff from the stage game, we have to describe how players evaluate infinite streams of payoffs $(u_i(a^1), u_i(a^2), \dots)$. Although there are alternative specifications in the literature, we will concentrate on the case of discounting, where players discount the period payoffs using a discount factor $\delta \in (0, 1)$. The payoff of player i to the infinite sequence $(a^1, a^2, \dots)$ is given by

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t).$$

The normalization factor $(1 - \delta)$ serves to measure the stage game and the repeated game payoffs in the same units. For example, the payoff to perpetual cooperation is given by

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \times 2 = 2.$$

The payoff of player i induced by strategy profile s is given by

$$U(s) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(a^t)$$

where a^t is the period t action profile if players comply with the strategy profile s . For example, if both players play according to the grim-trigger strategy profile, then period t action profile will be $a^t = (C, C)$ for all $t = 1, 2, \dots$, and hence

$$\begin{aligned} U(s) &= (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(C, C) \\ &= (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \times 2 \\ &= 2. \end{aligned}$$

Notice that each history starts a new subgame, and hence for any strategy profile s and history h^t , we can compute the players' expected payoffs from period t onward. We call these the **continuation payoffs** and renormalize so that the continuation payoffs from period t on are measured in period t units:

$$U_i(s|h^t) = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} u_i(a^{\tau+1})$$

if the strategy profile s induces the sequence of actions $(a^{t+1}, a^{t+2}, \dots)$ starting from history h^t .

Let us denote the resulting infinitely repeated game by G_δ .

3 Equilibria of Infinitely Repeated Games

Definition 1 *The strategy profile s is a **Nash equilibrium** of the repeated game G_δ if for all $i \in N$*

$$U_i(s_i, s_{-i}) \geq U_i(s'_i, s_{-i}) \text{ for all } s'_i \in S_i.$$

Let us consider some of the Nash equilibria of the PD game.

(1) Defection no matter what. Playing D after every history is clearly a Nash equilibrium. This is because whatever you do, your opponent will play D , and the best response to D is D as well.

(2) Grim-trigger. Let us now check if the grim-trigger strategy profile is a Nash equilibrium. Suppose player 2 adopts the grim-trigger strategy. If player 1 also follows the grim-trigger strategy then the outcome will be cooperation every period

$$(C, C), (C, C), \dots, (C, C), (C, C), \dots$$

with the resulting payoff sequence

$$2, 2, 2, \dots, 2, 2, 2, \dots$$

whose average discounted value is

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_i(C, C) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \times 2 = 2.$$

Now, consider the best possible deviation for player 1. For such a deviation to be profitable it must result in a sequence of action profiles which has defection by some players in some period. This, in turn, implies that player 1 must be defecting at some period (since player 2 is following grim-trigger she will not defect unless player 1 defected in the past). Let $T + 1$, where $T \in \{0, 1, \dots\}$, be the first period in which player 1 defects. Therefore, we have the following sequence of action profiles until period $T + 1$

$$\underbrace{(C, C), (C, C), \dots, (C, C)}_{T \text{ times}}, \underbrace{(D, C)}_{\text{period } T+1}.$$

Since player 2 is following the grim-trigger strategy she will play D in period $T + 2$ and after. Well, the best thing that player 1 can do in that case is to play D as well starting from period $T + 2$. Therefore, the best deviation by player 1 generates the following sequence of action profiles

$$\underbrace{(C, C), (C, C), \dots, (C, C)}_{T \text{ times}}, \underbrace{(D, C)}_{\text{period } T+1}, (D, D), (D, D), \dots$$

and the following sequence of period payoffs

$$\underbrace{2, 2, \dots, 2}_{T \text{ times}}, 3, 1, 1, \dots$$

The average discounted value of this sequence is

$$\begin{aligned} & (1 - \delta) [2 + \delta 2 + \delta^2 2 + \dots + \delta^{T-1} 2 + \delta^T 3 + \delta^{T+1} + \delta^{T+2} + \dots] \\ &= 2 + \delta^T - 2\delta^{T+1}. \end{aligned}$$

You can check to see that if $\delta \geq 1/2$, this is smaller than or equal to 2. Therefore, if players are patient enough, i.e., if $\delta \geq 1/2$, then grim-trigger strategy profile is a Nash equilibrium of the infinitely repeated PD game.

Definition 2 *The strategy profile s is a **subgame perfect equilibrium** of the repeated game G_δ if for all $i \in N$ and all $h^t \in H$*

$$U_i(s_i, s_{-i}|h^t) \geq U_i(s'_i, s_{-i}|h^t) \text{ for all } s'_i \in S_i.$$

Proposition 3 (One-Shot Deviation Property) *A strategy profile s^* is a SPE of G_δ if and only if no player can gain by changing her action after any history, keeping both the strategies of the other players and the remainder of her own strategy constant.*

Now that we have the one-deviation property at hand, we may analyze the SPE of the repeated prisoners' dilemma game.

Let us first consider a finitely repeated version. By backward induction it is easy to see that the only SPE in this case is defection (D) every period. This is because, D is strictly dominant in the last period T and hence both players play D after any history h^T . Now, in period $T - 1$ neither player can gain in period T by cooperating, and they lose in period $T - 1$, so that play in $T - 1$ will be defection as well after any history h^{T-1} . Continuing in this manner we have that both players will play D after every history h^t , $t = 1, 2, \dots, T$, in the unique SPE. [It turns out that this also the unique Nash equilibrium *outcome*. Prove this as an **exercise**].

Let us now consider the infinitely repeated version.

(1) Defection no matter what. It can be easily verified that the strategy profile given by

$$s_i(h^t) = D, \text{ for all } t = 1, 2, \dots \text{ and all } h^t,$$

for $i = 1, 2$, is subgame perfect (check using the one-shot deviation property).

(2) Grim-trigger. Now, let us consider the grim-trigger strategy. We have to check whether the grim trigger strategy satisfies the one-shot deviation property after every possible

history. Consider the history $h^2 = (C, D)$, i.e., in the first period player 2 defected. Let's see if player 2 has a profitable one-shot deviation. If player 2 plays according to the grim-trigger strategy, given that player 1 sticks to that strategy as well, the following sequence of action profiles will result (starting from period 2)

$$(D, C), (D, D), (D, D), \dots$$

with the sequence of payoffs

$$0, 1, 1, \dots$$

whose average discounted value is δ . If she deviates and plays D in the second period (after history (C, D)), keeping rest of her strategy the same, she will get a payoff of 1 every period, which has an average discounted value of 1. Therefore, this is a profitable deviation since $\delta < 1$, and the grim-trigger strategy profile is not a subgame perfect equilibrium.

(3) Grim-trigger II. We may, however, modify the grim-trigger strategy slightly and obtain a SPE. This strategy profile, which we will call grim-trigger II, is given by

$$s_i^*(h^t) = \begin{cases} C, & t = 1 \\ C, & h^t = ((C, C), (C, C), \dots, (C, C)) \\ D, & \text{otherwise} \end{cases}$$

for $i = 1, 2$. The difference in this strategy is that, a player defects if there is a defection in the past independent of the identity of the defector. We claim that s^* is a SPE for all $\delta \geq 1/2$.

Consider all histories of the type $h^t = ((C, C), (C, C), \dots, (C, C))$, i.e., all histories without any defection. For player 1, the conditional payoff to playing s_1^* is

$$U_1(s^*|h^t) = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \times 2 = 2.$$

One-shot deviation to D at period t gives

$$\begin{aligned} (1 - \delta) [u_1(D, C) + \delta + \delta^2 + \dots] &= (1 - \delta) 3 + \delta \\ &= 3 - 2\delta \\ &\leq 2 \end{aligned}$$

for all $\delta \geq 1/2$. Similarly, let h^t be a history other than $((C, C), (C, C), \dots, (C, C))$. Then,

$$U_1(s^*|h^t) = (1 - \delta) \sum_{\tau=t}^{\infty} \delta^{\tau-t} \times 1 = 1$$

whereas deviating and playing C in period t gives

$$\begin{aligned} (1 - \delta) [u_1(C, D) + \delta + \delta^2 + \dots] &= (1 - \delta) \times (0) + \delta \times 1 \\ &= \delta \\ &< 1. \end{aligned}$$

Similar considerations for player 2 shows that s^* is a SPE.

The grim-trigger strategies are very fierce in their punishments, they never forgive. We will now demonstrate that more forgiving strategies can sustain the cooperative outcome in a SPE as well.

(4) Forgiving trigger. In a **forgiving trigger strategy** a player starts with C and she plays C as long as everybody has played C in this past. If anybody plays D at any period, then the player plays D for the next k periods. After k periods of punishment she returns to playing C until someone deviates again.

Suppose the game is in the cooperative phase, i.e., either nobody has deviated so far or the deviations that have occurred have already been punished. We have to check whether there exists one-shot profitable deviation in this phase. Suppose player 2 follows the forgiving trigger strategy. If player 1 follows that strategy as well the outcome will be (C, C) forever which gives an average discounted payoff of 2. If player 1 deviates to D once and then follows the strategy, then the following sequence of actions will result

$$(D, C), \underbrace{(D, D), (D, D), \dots, (D, D)}_{k \text{ times}}, (C, C), (C, C), \dots$$

with the following average discounted payoff

$$(1 - \delta) [3 + \delta + \delta^2 + \dots + \delta^k + 2\delta^{k+1} + 2\delta^{k+2} + \dots] = 3 - 2\delta + \delta^{k+1}.$$

Therefore, there is no profitable one-shot deviation in the cooperative phase if and only if

$$3 - 2\delta + \delta^{k+1} \leq 2$$

or

$$\delta^{k+1} - 2\delta + 1 \leq 0.$$

If, for example, $k = 1$, then this condition is equivalent to

$$\delta^2 - 2\delta + 1 = (\delta - 1)^2 \leq 0$$

which can never be satisfied since $\delta < 1$. If, however, $k = 2$, then the condition will be satisfied for any $\delta \geq 0.62$. In general, as the length of the punishment phase increases, the lower bound on δ decreases and converges to $1/2$ as $k \rightarrow \infty$.

We also have to check if there is a profitable one-shot deviation in the punishment phase. Suppose there are $k' \leq k$ periods left in the punishment phase. If player 1 complies with the forgiving trigger strategy the following action profiles will be realized

$$\underbrace{(D, D), (D, D), \dots, (D, D)}_{k' \text{ times}}, (C, C), (C, C), \dots$$

and if she deviates once at the beginning

$$\underbrace{(C, D), (D, D), \dots, (D, D)}_{k' \text{ times}}, (C, C), (C, C), \dots$$

Clearly, following the forgiving trigger strategy is better in the punishment phase.

This establishes that if $k \geq 2$ and the players are patient enough, then forgiving trigger strategies form a SPE.

So far we have analyzed only the PD game to illustrate some of the results that can be obtained in repeated games. The repeated games literature considers all possible games and characterizes the set of possible outcomes that can be obtained in the Nash equilibria or SPE of repeated games. Results, known as “folk theorems”, have shown that virtually any outcome can be supported as a Nash and SPE outcome in infinitely repeated games, provided that the players are patient enough. (see Fudenberg and Tirole, *Game Theory*, 1991 for more on this.)

We will end our discussion of repeated games with another example, this time from industrial-organization theory.

Example 4 Consider a Cournot duopoly model with inverse demand function

$$P(Q) = \begin{cases} a - Q, & Q \leq a \\ 0, & Q > a \end{cases}$$

where $Q = Q_1 + Q_2$ and cost functions $C_i(Q_i) = cQ_i$, $i = 1, 2$. The profit function of firm i is given by

$$u_i(Q_1, Q_2) = Q_i(P(Q_1 + Q_2) - c).$$

Consider the following grim-trigger strategy. Produce half the monopoly output in the first period and as long as everybody has produced that amount so far. Otherwise produce the Cournot output. As an exercise verify that this is a SPE.

Extensive Form Games with Incomplete Information

Levent Koçkesen

1 Introduction

So far we have analyzed games in strategic form with and without incomplete information, and extensive form games with complete information. In this section we will analyze extensive form games with incomplete information. Many interesting strategic interactions can be modelled in this form, such as signalling games, repeated games with incomplete information in which reputation building becomes a concern, bargaining games with incomplete information, etc.

The analysis of extensive form games with incomplete information will show that we need to provide further refinements of the Nash equilibrium concept. In particular, we will see that subgame perfect equilibrium (SPE) concept that we have introduced when we studied extensive form games with complete information is not adequate. To illustrate the main problem in the SPE concept, however, the following game with imperfect, but complete, information is sufficient.




Figure 1: Something Wrong with SPE

The strategic form of this game is given by

	<i>L</i>	<i>R</i>
<i>O</i>	1, 3	1, 3
<i>T</i>	2, 1	0, 0
<i>B</i>	0, 2	0, 1

It can be easily seen that the set of Nash equilibria of this game is $\{(T, L), (O, R)\}$. Since this game has only one subgame, i.e., the game itself, this is also the set of SPE. But there

is something implausible about the (O, R) equilibrium. Action R is strictly dominated for player 2 at the information set I . Therefore, if the game ever reaches that information set, player 2 should never play R . Knowing that, then, player 1 should play T , as she would know that player 2 would play L , and she would get a payoff of 2 which is bigger than the payoff that she gets by playing O . Subgame perfect equilibrium cannot capture this, because it does not test rationality of player 2 at the non-singleton information set I .

The above discussion suggests the direction in which we have to strengthen the SPE concept. We would like players to be rational not only in every subgame but also in every **continuation game**. A continuation game in the above example is composed of the information set I and the nodes that follow from that information set. First, notice that the continuation game does not start with a single decision node, and hence it is not a subgame. However, rationality of player 2 requires that he plays action L if the game ever reaches there.

In general, the optimal action at an information set may depend on which node in the information set the play has reached. Consider the following modification of the above game.




Figure 2:

Here the optimal action of player 2 at the information set I depends on whether player 1 has played T or B - information that 2 does not have. Therefore, analyzing player 2's decision problem at that information set requires him forming beliefs regarding which decision node he is at. In other words, we require that

(1) (Condition 1: Beliefs) At each information set the player who moves at that information set has beliefs over the set of nodes in that information set.

and

(2) (Condition 2: Sequential Rationality) At each information set, strategies must be optimal, given the beliefs and subsequent strategies.

Let us check what these two conditions imply in the game given in Figure 1. Condition 1 requires that player 2 assigns beliefs to the two decision nodes at the information set I . Let

the probability assigned to the node that follows T be $\mu \in [0, 1]$ and the one assigned to the node that follows B be $1 - \mu$. Given these beliefs the expected payoff to action L is

$$\mu \times 1 + (1 - \mu) \times 2 = 2 - \mu$$

whereas the expected payoff to R is

$$\mu \times 0 + (1 - \mu) \times 1 = 1 - \mu.$$

Notice that $2 - \mu > 1 - \mu$ for any $\mu \in [0, 1]$. Therefore, Condition 2 requires that in equilibrium player 2 never plays R with positive probability. This eliminates the subgame perfect equilibrium (O, R) , which, we argued, was implausible.

Although it requires players to form beliefs at non-singleton information sets, condition 1, does not specify how these beliefs are formed. As we are after an equilibrium concept, we require the beliefs to be consistent with the players' strategies. As an example consider the game given in Figure 2 again. Suppose player 1 plays actions O , T , and B with probabilities $\beta_1(O)$, $\beta_1(T)$, and $\beta_1(B)$, respectively. Also let μ be the belief assigned to node that follows T in the information set I . If, for example, $\beta_1(T) = 1$ and $\mu = 0$, then we have a clear inconsistency between player 1's strategy and player 2's beliefs. The only consistent belief in this case would be $\mu = 1$. In general, we may apply **Bayes' Rule**, whenever possible, to achieve consistency:

$$\mu = \frac{\beta_1(T)}{\beta_1(T) + \beta_1(B)}.$$

Of course, this requires that $\beta_1(T) + \beta_1(B) \neq 0$. If $\beta_1(T) + \beta_1(B) = 0$, i.e., player 1 plays action O with probability 1, then player 2 does not obtain any information regarding which one of his decision nodes has been reached from the fact that the play has reached I . The weakest consistency condition that we can impose is then,

(3) (Condition 3: Weak Consistency) Beliefs are determined by Bayes' Rule and strategies whenever possible.

These three conditions define the equilibrium concept **Perfect Bayesian Equilibrium** (PBE).

2 Perfect Bayesian Equilibrium

To be able to define PBE more formally, let H_i be the set of all information sets a player has in the game, and let $A(h)$ be the set of actions available at information set h . A **behavioral strategy** for player i is a function β_i which assigns to each information set $h \in H_i$ a probability distribution on $A(h)$, i.e.,

$$\sum_{a \in A(h)} \beta_i(a) = 1.$$

Let $\mathcal{B}_i$ be the set of all behavioral strategies available for player i and $\mathcal{B}$ be the set of all behavioral strategy profiles, i.e., $\mathcal{B} = \times_i \mathcal{B}_i$. A **belief system** $\mu : X \rightarrow [0, 1]$ assigns to each decision node x in the information set h a probability $\mu(x)$, where $\sum_{x \in h} \mu(x) = 1$ for all $h \in H$. Let M be the set of all belief systems. An **assessment** $(\mu, \beta) \in M \times \mathcal{B}$ is a belief system combined with a behavioral strategy profile.

Perfect Bayesian equilibrium is an assessment (μ, β) that satisfy conditions 1-3.¹ To illustrate, consider the game in Figure 2 again. Let $\beta_i(a)$ be the probability assigned to action a by player i , and μ be the belief assigned to the node that follows T in information set I . In any PBE of this game we have (i) $\beta_2(L) = 1$, (ii) $\beta_2(L) = 0$, or (iii) $\beta_2(L) \in (0, 1)$. Let us check each of the possibilities in turn:

(i) $\beta_2(L) = 1$. In this case, sequential rationality of player 2 implies that the expected payoff to L is greater than or equal to the expected payoff to R , i.e.,

$$\mu \times 1 + (1 - \mu) \times 1 \geq \mu \times 0 + (1 - \mu) \times 2$$

or

$$1 \geq 2 - 2\mu \iff \mu \geq 1/2.$$

Sequential rationality of player 1 on the other hand implies that she plays T , i.e., $\beta_1(T) = 1$. Bayes' rule then implies that

$$\mu = \frac{\beta_1(T)}{\beta_1(T) + \beta_1(B)} = \frac{1}{1 + 0} = 1,$$

which is greater than $1/2$, and hence does not contradict player 2's sequential rationality. Therefore, the following assessment is a PBE

$$\beta_1(T) = 1, \beta_2(L) = 1, \mu = 1.$$

(ii) $\beta_2(L) = 0$. Sequential rationality of player 2 now implies that $\mu \leq 1/2$, and sequential rationality of player 1 implies that $\beta_1(O) = 1$. Since $\beta_1(T) + \beta_1(B) = 0$, however, we cannot apply Bayes' rule, and hence condition 3 is trivially satisfied. Therefore, there is a continuum of equilibria of the form

$$\beta_1(O) = 1, \beta_2(L) = 0, \mu \leq 1/2.$$

(iii) $\beta_2(L) \in (0, 1)$. Sequential rationality of player 2 implies that $\mu = 1/2$. For player 1 the expected payoff to O is 1, to T is $2\beta_2(L)$, and to B is 0. Clearly, player 1 will never play B

¹Perfect Bayesian equilibrium, as it was defined in D. Fudenberg and J. Tirole (1991), "Perfect Bayesian and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236-60, puts more restrictions on the out-of-equilibrium beliefs and hence is stronger than the definition provided here.

with positive probability, that is in this case we always have $\beta_1(B) = 0$. If, $\beta_1(O) = 1$, then we must have $2\beta_2(L) \leq 1 \iff \beta_2(L) \leq 1/2$, and we cannot apply Bayes' rule. Therefore, any assessment that has

$$\beta_1(O) = 1, 0 < \beta_2(L) \leq 1/2, \mu = 1/2$$

is a PBE. If, on the other hand, $\beta_1(O) = 0$, then we must have $\beta_1(T) = 1$, and Bayes' rule implies that $\mu = 1$, contradicting $\mu = 1/2$. If $\beta_1(O) \in (0, 1)$, then Bayes' rule implies that $\mu = 1$, again contradicting $\mu = 1/2$.

Perfect Bayesian Equilibrium could be considered a weak equilibrium concept, because it does not put enough restrictions on out-of-equilibrium beliefs. Consider the three-player game given in Figure 3. The unique subgame perfect equilibrium of this game is (D, L, R') . However, the strategy profile (A, L, L') together with the belief system that puts probability 1 to the node that follows R is an assessment that satisfies conditions 1-3. Clearly, this is not a plausible outcome, as (L, L') is not a Nash equilibrium of the subgame that starts with player 2's move. Also, notice that player 3's beliefs are not consistent with player 2's strategy, but since player 3's information set is off-the-equilibrium, Bayes' rule has no bite there.




Figure 3: PBE may have "unreasonable" beliefs

The most commonly used equilibrium concept that do not suffer from such deficiencies is that of sequential equilibrium. Before we can define sequential equilibrium, however, we have to define a particular consistency notion. A behavioral strategy profile is said to be **completely mixed** if every action receives positive probability.

Definition (Consistency): An assessment (μ, β) is **consistent** if there exists a completely mixed sequence (μ^n, β^n) that converges to (μ, β) such that μ^n is derived from β^n using Bayes' rule for all n .

An assessment (μ, β) is a **sequential equilibrium** if it is sequentially rational and consistent. To illustrate, consider the game in Figure 3 again. Let μ be the probability assigned to the node that follows L , and consider the assessment $((A, L, L'), \mu = 0)$. For this to be a sequential equilibrium, we have to find a completely mixed behavioral strategy profile β^n such that

$$\beta_1^n(A) \rightarrow 1, \beta_2^n(L) \rightarrow 1, \beta_3^n(L') \rightarrow 1, \mu^n = \frac{\beta_2^n(L)}{\beta_2^n(L) + \beta_2^n(R)} \rightarrow 0,$$

which is not possible. However, the assessment given by $((D, L, R'), \mu = 1)$ is easily checked to satisfy sequential rationality. To check consistency, let

$$\beta_1^n(D) = 1 - \frac{1}{n}, \beta_2^n(L) = 1 - \frac{1}{n}, \beta_3^n(R') = 1 - \frac{1}{n}, \mu^n = 1 - \frac{1}{n}.$$

Notice that μ^n is derived from β^n via Bayes' rule and $(\mu^n, \beta^n) \rightarrow (\mu, \beta)$. Therefore, this assessment is a sequential equilibrium.

3 Signalling Games

One of the most common applications in economics of extensive form games with incomplete information is signalling games. In its simplest form, in a signalling game there are two players, a sender S , and a receiver, R . Nature draws the type of the sender from a type set Θ , whose typical element will be denoted θ . The probability of type θ being drawn is $p(\theta)$. Sender observes his type and chooses a message $m \in M$. The receiver observes m (but not θ) and chooses an action $a \in A$. The payoffs are given by $u_S(m, a, \theta)$ and $u_R(m, a, \theta)$.

Let $\mu(\theta|m)$ denote the receiver's belief that the sender's type is θ if message m is observed. Also let $\beta_S(m|\theta)$ denote the probability that type θ sender sends message m , and $\beta_R(a|m)$ denote the probability that the receiver chooses action a after observing message m . Given an assessment (μ, β) , the expected payoff of a sender of type θ is then

$$U_S(\mu, \beta, \theta) = \sum_m \sum_a \beta_S(m|\theta) \beta_R(a|m) u_S(m, a, \theta),$$

whereas the expected payoff of the receiver conditional upon receiving message m is

$$U_R(\mu, \beta|m) = \sum_\theta \sum_a \mu(\theta|m) \beta_R(a|m) u_R(m, a, \theta).$$

Also, Bayes' rule implies,

$$\mu(\theta'|m') = \frac{\beta_S(m'|\theta') p(\theta')}{\sum_\theta \beta_S(m'|\theta) p(\theta)},$$

whenever $\sum_\theta \beta_S(m'|\theta) p(\theta) \neq 0$, i.e., at least one type of sender sends the message m' .

To illustrate consider the game in figure 4, known as **Beer or Quiche**. In this game Nature (N) chooses the type of player 1 to be Tough (T) (with probability 0.9) or Weak (W)

(with probability 0.1). Player 1 observes her type and chooses Quiche (Q) or Beer (B). Player 2 observes only the action choice of player 1 but not the type, and chooses to fight (F) or not to fight (A).




Figure 4: Beer or Quiche

Let us find the pure strategy PBE of this game. There are four types of possible equilibria:

1. Each type chooses a different action (**Separating Equilibria**):

- (a) Weak chooses quiche, Tough chooses beer ($\beta_S(Q|W) = 1, \beta_S(Q|T) = 0$) :

Bayes' rule implies that

$$\mu(W|Q) = \frac{\beta_S(Q|W)p(W)}{\beta_S(Q|W)p(W) + \beta_S(Q|T)p(T)} = \frac{1 \times 0.1}{1 \times 0.1 + 0 \times 0.9} = 1.$$

Similarly, $\mu(T|B) = 1$. Therefore, the receiver's sequential rationality implies that $\beta_R(A|B) = 1$ and $\beta_R(F|Q) = 1$. Sequential rationality of the sender, then, implies that $\beta_S(Q|W) = 0$, contradicting our hypothesis. So, there is no PBE of this type.

- (b) Weak chooses beer, Tough chooses quiche ($\beta_S(Q|W) = 0, \beta_S(Q|T) = 1$) :

Bayes' rule implies that $\mu(T|Q) = 1$ and $\mu(W|B) = 1$. Therefore, the receiver's sequential rationality implies that $\beta_R(F|B) = 1$ and $\beta_R(A|Q) = 1$. Sequential rationality of the sender, then, implies that $\beta_S(Q|W) = 1$, contradicting our hypothesis. So, there is no PBE of this type either.

2. Both types choose the same action (**Pooling Equilibria**)

(a) Both choose quiche ($\beta_S(Q|W) = 1, \beta_S(Q|T) = 1$) :

Bayes' rule implies that $\mu(W|Q) = 0.1$ and $\mu(T|Q) = 0.9$. Therefore, after observing Q , the receiver's expected payoff to F is

$$0.1 \times 1 + 0.9 \times 0 = 0.1$$

and expected payoff to A is

$$0.1 \times 0 + 0.9 \times 2 = 1.8,$$

and hence sequential rationality implies that $\beta_R(A|Q) = 1$. The weak type's sequential rationality implies, then, that $\beta_S(Q|W) = 1$, confirming our hypothesis. For the tough type, playing quiche would be rational only if the receiver chooses to fight after observing beer. Therefore, we must have $\beta_R(F|B) = 1$, which in turn requires that $\mu(W|B) \geq 1/2$. Therefore, any assessment which satisfies the following is a PBE:

$$\begin{aligned}\beta_S(Q|W) &= 1, \beta_S(Q|T) = 1, \beta_R(A|Q) = 1, \beta_R(F|B) = 1, \\ \mu(W|Q) &= 0.1, \mu(W|B) \geq 1/2.\end{aligned}$$

(b) Both choose beer ($\beta_S(B|W) = 1, \beta_S(B|T) = 1$) : It is easily checked that the following constitute the set of PBE of this type:

$$\begin{aligned}\beta_S(B|W) &= 1, \beta_S(B|T) = 1, \beta_R(F|Q) = 1, \beta_R(A|B) = 1, \\ \mu(W|B) &= 0.1, \mu(W|Q) \geq 1/2.\end{aligned}$$

Job Market Signalling²

Suppose there are two types of workers, a high ability (H) and a low ability (L) type. We let the probability of having high ability be denoted by $p \in (0, 1)$. The output is equal to 2 if the worker is of high ability and equal to 1 if he is of low ability. The worker can choose a level of education $e \geq 0$ before applying for a job. However, the cost of having level of education e is e for the low ability worker and $e/2$ for the high ability worker. The worker knows his ability but the employer observes only the level of education, not the ability. Therefore, the employer offers a wage schedule $w(e)$ as a function of education. The payoffs of the workers are given by

$$\begin{aligned}u(w, e, H) &= w - e/2, \\ u(w, e, L) &= w - e.\end{aligned}$$

²Based on M. Spence (1973), "Job Market Signalling," *Quarterly Journal of Economics*, 87, 355-74.

We assume that the job market is competitive and hence the employer offers a wage schedule $w(e)$ such that the expected profit is equal zero. Therefore, if $\mu(H|e)$ denotes the belief of the employer that the worker is of high ability given that he has chosen education level e , the wage schedule will satisfy $w(e) = 2\mu(H|e) + (1 - \mu(H|e))$. We are interested in the set of PBE of this game. Let e_H and e_L denote the education levels chosen by the high and low ability workers, respectively.

1. **Separating Equilibria** ($e_H \neq e_L$): The Bayes' rule in this case implies that $\mu(H|e_H) = 1$ and $\mu(L|e_L) = 1$. Therefore, we have $w(e_H) = 2$ and $w(e_L) = 1$. Given that, the low ability worker will choose $e = 0$. In equilibrium, it must be such that the low ability worker does not want to mimic the high ability worker and vice versa. Therefore, we need to have

$$2 - \frac{e_H}{2} \geq 1$$

or $e_H \leq 2$ and

$$1 \geq 2 - e_H$$

or $1 \leq e_H$. We can support any e_H between 1 and 2 with the following belief system

$$\mu(H|e) = \begin{cases} 0, & e < e_H \\ 1, & e \geq e_H \end{cases}.$$

2. **Pooling Equilibria** ($e_H = e_L = e^*$): The Bayes' rule in this case implies that $\mu(H|e^*) = p$ and $\mu(L|e^*) = 1 - p$. Therefore, $w(e^*) = 2p + (1 - p) = p + 1$ and hence

$$\begin{aligned} u(w, e^*, H) &= p + 1 - e^*/2, \\ u(w, e^*, L) &= p + 1 - e^*. \end{aligned}$$

It must be the case that

$$\begin{aligned} p + 1 - e^*/2 &\geq 0 \\ p + 1 - e^* &\geq 0. \end{aligned}$$

We also need to have

$$\begin{aligned} p + 1 - e^*/2 &\geq w(e) - e/2, \\ p + 1 - e^* &\geq w(e) - e, \end{aligned}$$

for all $e \geq 0$. The above inequalities are satisfied if and only if $e^* \leq p$. We can, in turn, show that any such e^* can be supported as an equilibrium by the following belief system

$$\mu(H|e) = \begin{cases} p, & e = e^* \\ 0, & e \neq e^* \end{cases}$$